



NAVAL POSTGRADUATE SCHOOL

MONTEREY, CALIFORNIA

THESIS

**AN ANALYSIS OF MATHEMATICAL MODELS TO
IMPROVE COUNTER-DRUG SMUGGLING
OPERATIONS**

by

Sylvester Campos III

September 2014

Thesis Advisor:
Second Reader:

Michael Atkinson
Johannes O. Royset

Approved for public release; distribution is unlimited

THIS PAGE INTENTIONALLY LEFT BLANK

REPORT DOCUMENTATION PAGE			<i>Form Approved OMB No. 0704-0188</i>	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instruction, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188) Washington DC 20503.				
1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE September 2014	3. REPORT TYPE AND DATES COVERED Master's Thesis	
4. TITLE AND SUBTITLE AN ANALYSIS OF MATHEMATICAL MODELS TO IMPROVE COUNTER-DRUG SMUGGLING OPERATIONS			5. FUNDING NUMBERS	
6. AUTHOR(S) Sylvester Campos III				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Postgraduate School Monterey, CA 93943-5000			8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING /MONITORING AGENCY NAME(S) AND ADDRESS(ES) N/A			10. SPONSORING/MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES The views expressed in this thesis are those of the author and do not reflect the official policy or position of the Department of Defense or the U.S. Government. IRB protocol number ____N/A____.				
12a. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release; distribution is unlimited			12b. DISTRIBUTION CODE A	
13. ABSTRACT (maximum 200 words) <p>The Joint Interagency Task Force South (JIATF South) serves as the center for counterillicit trafficking operations against Central and South American drug cartels. JIATF South has jurisdiction over an enormous area and has a limited number of search and interdiction assets. Thus, mathematical models and tools to improve the search planning process could provide an immense benefit. This thesis contributes to a multiinstitution effort to develop these models and tools. The goal is to provide enhanced situational awareness to operators about the likely location of smuggling targets and recommend search plans.</p> <p>The thesis combines a probability model and an optimization model into a user-friendly tool that operators and researchers can easily use. We also develop methods to convert the output of the probability model into inputs to the optimization model. This could be especially important if environmental factors play a significant role in the routes taken by drug smugglers. We examine how accurately this conversion algorithm reproduces the actual paths taken by smugglers. More importantly, we evaluate how well the search plan generated by the conversion algorithm performs. We find, in most cases, that the conversion algorithm provides excellent results; however, there are some situations where the conversion algorithm performs poorly.</p>				
14. SUBJECT TERMS probability, optimization, multivariate adaptive regression splines (MARS), counter-drug smuggling operations, transnational criminal organizations (TCO), Joint Interagency Task Force South (JIATF South), drug smugglers, illicit trade			15. NUMBER OF PAGES 95	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT UU	

THIS PAGE INTENTIONALLY LEFT BLANK

Approved for public release; distribution is unlimited

**AN ANALYSIS OF MATHEMATICAL MODELS TO IMPROVE COUNTER-
DRUG SMUGGLING OPERATIONS**

Sylvester Campos III
Lieutenant, United States Navy
B.S., University of North Florida, 2008

Submitted in partial fulfillment of the
requirements for the degree of

MASTER OF SCIENCE IN OPERATIONS RESEARCH

from the

**NAVAL POSTGRADUATE SCHOOL
September 2014**

Author: Sylvester Campos III

Approved by: Michael Atkinson
Thesis Advisor

Johannes O. Royset
Second Reader

Robert F. Dell
Chair, Department of Operations Research

THIS PAGE INTENTIONALLY LEFT BLANK

ABSTRACT

The Joint Interagency Task Force South (JIATF South) serves as the center for counterillicit trafficking operations against Central and South American drug cartels. JIATF South has jurisdiction over an enormous area and has a limited number of search and interdiction assets. Thus, mathematical models and tools to improve the search planning process could provide an immense benefit. This thesis contributes to a multiinstitution effort to develop these models and tools. The goal is to provide enhanced situational awareness to operators about the likely location of smuggling targets and recommend search plans.

The thesis combines a probability model and an optimization model into a user-friendly tool that operators and researchers can easily use. We also develop methods to convert the output of the probability model into inputs to the optimization model. This could be especially important if environmental factors play a significant role in the routes taken by drug smugglers. We examine how accurately this conversion algorithm reproduces the actual paths taken by smugglers. More importantly, we evaluate how well the search plan generated by the conversion algorithm performs. We find, in most cases, that the conversion algorithm provides excellent results; however, there are some situations where the conversion algorithm performs poorly.

THIS PAGE INTENTIONALLY LEFT BLANK

TABLE OF CONTENTS

I.	INTRODUCTION.....	1
A.	BACKGROUND	1
B.	MARITIME TRANSNATIONAL CRIMINAL ORGANIZATIONS	4
C.	UNITED STATES COUNTERDRUG ORGANIZATIONS	6
1.	Joint Interagency Task Force South	6
2.	United States Coast Guard.....	7
3.	United States Navy	7
D.	LAW OF THE SEA	8
E.	IMPACT OF ILLICIT DRUGS ON THE UNITED STATES.....	9
F.	MODELING APPROACH	10
G.	THESIS STRUCTURE	12
II.	LITERATURE REVIEW	15
A.	SEARCH THEORY.....	15
B.	SMUGGLER SEARCH PROBLEM	16
C.	PROBABILITY MODELS	18
D.	MULTIVARIATE ADAPTIVE REGRESSION SPLINES.....	19
III.	DEVELOPMENT OF THE TOOL.....	23
A.	GENERATING PROBABILITY HEAT MAPS	23
B.	OPTIMIZATION MODEL	27
1.	Generating Search Plans Using the Optimization Model	28
2.	Options File.....	29
3.	Search and Detection Capabilities.....	29
4.	JIATF South Search Assets	31
5.	Case Information	32
C.	OPTIMIZATION OUTPUT	32
D.	OPTIMIZED SEARCH PLAN OUTPUT	33
E.	CONVERSION ALGORITHM.....	35
IV.	GENERATING OPTIMIZATION INPUT FROM HEAT MAPS	37
A.	TIMES.....	38
B.	GENERAL PATHFINDING MODEL APPROACH	39
1.	Expected Track	39
2.	MARS Function	41
3.	Area of Uncertainty	43
4.	Determining Waypoint Location	46
V.	RESULTS	49
A.	CASE GENERATION.....	49
B.	SINGLE TARGET RESULTS	50
1.	Case 1: Straight Path Route with Changing Velocity.....	50
2.	Case 2: Multiple Leg Route.....	54
3.	Case 3: Zig-Zag Route	57

4.	Quick Change of Heading	60
C.	MULTIPLE TARGET RESULTS	62
D.	COMPARISON ACROSS CASES.....	65
VI.	SUMMARY AND FUTURE WORK	67
A.	SUMMARY	67
B.	FUTURE WORK	68
	LIST OF REFERENCES	71
	INITIAL DISTRIBUTION LIST	75

LIST OF FIGURES

Figure 1.	Self-Propelled Semisubmersible (from <i>On Patrol Magazine</i> , 2013).....	2
Figure 2.	Movement of Cocaine into the United States (from Customs Border Protection [CBP], 2014).....	3
Figure 3.	Panga Smuggling Narcotics off the Coast of California (from Military.com, 2014).	6
Figure 4.	Legal Boundaries of Oceans and Airspace (from United Nations, 1982).	9
Figure 5.	Snapshot of a Probability Heat Map.	11
Figure 6.	Example of a PARS “Heat Map” Output (from Sidoti et al., 2013).	18
Figure 7.	Sample Counterdrug Probability Heat Map.	20
Figure 8.	Illustration of MARS vs. OLS Longitude vs. Time for a Target.	21
Figure 9.	Heat Map Representation of a Target’s Path.	25
Figure 10.	Depiction of the Area of Interest and a Single Target Scenario.	27
Figure 11.	All Optimization Model Search Assets.....	32
Figure 12.	Illustration of the Search Plan.....	34
Figure 13.	Snapshot of Probability Heat Map at One Time Step.....	40
Figure 14.	A Target’s Track through Time.	40
Figure 15.	Expected Track of Longitude vs. Latitude.....	41
Figure 16.	Illustrates the MARS Model Fit (Red) vs. the Linear Regression Model Fit (Blue).	42
Figure 17.	Capturing Changes in Heading or Velocity Using MARS.	43
Figure 18.	Calculating Area of Uncertainty for a Single Time Step.....	46
Figure 19.	Sample Pathfinding Model Output.	47
Figure 20.	The Actual Tracks of Eight Cases of Interest.	50
Figure 21.	Snapshot of Area of Uncertainty for Case 1.	51
Figure 22.	Search Plan Results for Case 1.	52
Figure 23.	Case 1 MARS Model Fit.....	53
Figure 24.	Snapshots of the Heat Map for Case 2.....	55
Figure 25.	Search Plan Results for the Multiple Leg Case.	55
Figure 26.	Single Target, Multiple Leg MARS Model Fit.....	56
Figure 27.	Snapshot of the Corner-Cutting Effect.	56
Figure 28.	Snapshot of the Area of Uncertainty for the Zig-Zag Case.	58
Figure 29.	Search Plan Results for the Zig-Zag Case.	58
Figure 30.	MARS Model Fit for Zig-Zag Case.....	59
Figure 31.	Snapshot of the AOU for Case 4.....	60
Figure 32.	Search Plan Results for Case 4.	61
Figure 33.	MARS Output for Case 4.....	62
Figure 34.	Snapshot of the Targets Area of Uncertainty for Case 5.	63
Figure 35.	Search Plan Results for Multitarget Scenario.	63
Figure 36.	MOP-2 Difference Between Time-On-Station Results.	65
Figure 37.	MOP-3 Percentage of Drugs Seized as a Result of the Search Plan.....	66

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF TABLES

Table 1.	Sample Data of the Case Information File.....	24
Table 2.	Sample Sweep-Width Data.	30
Table 3.	Sample Search Asset Information.....	31
Table 4.	All Search Assets Search for Target 1.	33
Table 5.	Sample Detailed Output CSV File for Searcher's Location over Time.....	33
Table 6.	Sample Data of the Case Information File.....	38

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF ACRONYMS AND ABBREVIATIONS

ADAM-II	Arrestee Drug Abuse Monitoring Program
AOI	area of interest
AOR	area of responsibility
AOU	area of uncertainty
CBP	U.S. Customs and Border Protection
CIA	Central Intelligence Agency
CSV	comma separated value
DEA	Drug Enforcement Administration
DHS	Department of Homeland Security
DOD	Department of Defense
DTO	Drug Trafficking Organization
FBI	Federal Bureau of Investigation
FFG	Oliver Hazard Perry-Class Frigate
FT	feet
GAMS	General Algebraic Modeling System
GOP-RDR	Generalized Orienteering Problem w/Resource Dependent Rewards
HR/HRS	hour/hours
ID	identification
ISR	intelligence, surveillance, and reconnaissance
JIATF-S	Joint Interagency Task Force-South
KG	kilogram
KM	kilometer
KP	knapsack problem
LCS	littoral combat ship
LEDET	law enforcement detachment
MARS	Multivariate Adaptive Regression Splines
METOC	meteorology and oceanography
MOP	measure of performance

NM	nautical miles
PARS	pirate attack risk surface
NDIC	National Drug Intelligence Center
NDTS	National Drug Threat Survey
NPS	Naval Postgraduate School
NRL	Naval Research Laboratory
OLS	ordinary least squares
OP	orienteering problem
RHIB	Rigid Hull Inflatable Boat
SOUTHCOM	United States Southern Command
SPAWAR	Space and Naval Warfare Systems Command
SPSS	self-propelled semisubmersible
SSP	smuggler search problem
TBM	theater ballistic missile
TCO	transnational criminal organization
TOS	time on station
TSP	Travelling Salesperson Problem
USCG	United States Coast Guard
USN	United States Navy
UCONN	University of Connecticut
UNCLOS	United Nations Convention on Law of the Sea

EXECUTIVE SUMMARY

This research is motivated by the ongoing efforts of the Joint Interagency Task Force South (JIATF South), which conducts search operations in order to stem the flow of illicit traffickers out of South and Central America. Planning real-world search operations is a particularly difficult task. Planners, operating under strict time constraints, weigh uncertain information about target whereabouts against the limitations of their available search assets. In an effort to assist JIATF South in combatting transnational criminal organizations, this thesis examines probability and optimization models for the maritime domain that will increase the effectiveness of counterdrug operations.

The probability model estimates the likely location of smugglers over time and the optimization model recommends routes for search assets to detect those smugglers. Prior to this thesis, the probability and optimization models were implemented in two separate pieces of code that were difficult for any non-expert to use. Thus, the first phase of this thesis is to develop a user friendly tool that combines the probability model with the optimization model. We want to create a buffer between the user and those technical details, so the user only has to enter in standard inputs, such as waypoints and velocities. The outputs to the algorithms are similarly difficult to parse for a nonexpert. Thus, we also create a function that will display the output in a more user-friendly form that novices will understand. This tool, with simplified input and output functionality, makes the models much more accessible.

Both the probability model and optimization model require intelligence information as input. The intelligence input consists of waypoint locations, departure times, velocities, and drug loads. The probability model produces a heat map that specifies the likely location of the targets. Environmental factors may influence how the heat map blobs move across the map. Currently, there is no connection between the probability model and the optimization model. A connection between the two models may not be necessary for the optimization model to produce operationally effective plans. However, the probability model can provide the optimization model with enhanced data about how environmental factors may influence a targets route. When the environment

significantly impacts how and where the drug smugglers transit, then the data being input into the optimization model may produce better search plans. Thus, the second part of the thesis develops a pathfinding algorithm that converts the heat map output of the probability model into the optimization model inputs.

In order to test how well our conversion algorithm performs, we compare the target's track given by the input intelligence to the track produced by the conversion algorithm. More importantly, we also examine the search plans produced. In most situations, our estimated track via the conversion algorithm corresponds closely to the actual track. In some situations, the conversion algorithm produces a "corner cutting" effect when the target changes heading. The algorithm correctly determines the waypoint, but the exact location of the waypoint is slightly off. The algorithm can also perform poorly when the target changes heading significantly after a short period of time. We do not have enough information in these cases to accurately determine the heading change.

Finally, we define three measures of performance to evaluate how well a search plan performs: total distance between searcher and target, total on-station time, and expected quantity of drugs seized due to search. In most cases, both the search plan generated from actual intelligence and the plan generated from the conversion algorithm perform well with regard to the distance between searcher and target. Not surprisingly, the conversion algorithm plan performed poorly in the cases where the estimated track did not match the actual track. In a few cases, however, where the tracks did not match that well, the search plan still did well.

With regard to time-on-station, the search plan based on actual intelligence outperformed the plan based on the conversion algorithm. On average, the searcher had one hour less search time for the conversion algorithm plan. This is not surprising, as the extra layer of processing introduced by the conversion algorithm distorts the information. Similarly, the actual search plan performed better with regard to expected drugs detected. The conversion algorithm plan only detected 75 percent as many drugs as the actual plan.

ACKNOWLEDGMENTS

First and foremost, I want to give praise to the Lord for being my light and my salvation. I would like to thank my wife, Danielle; her loving support gives me strength and allowed me to devote so much time to this thesis. I am ever indebted to my advisor, Professor Michael Atkinson, for his continued support and guidance throughout the arduous thesis process. I would like to thank my second reader, Professor Johannes Royset, for helping me through the thesis process. I would like to thank Dr. Jim Hansen, the developer of the PARS model, for creating the foundation that this thesis topic is based on. I would like to thank Dr. Jesse Pietz, the developer of the Generalized Orienteering Problem for Optimal Search and Interdiction Planning for Optimal Search and Interdiction, for giving me permission to continue his work in further developing this tool. Finally, I would like to thank the NPS Operations Research professors for their devotion and commitment to teaching me, the military officer, and the world of academia.

THIS PAGE INTENTIONALLY LEFT BLANK

I. INTRODUCTION

The Joint Interagency Task Force South (JIATF South) works alongside partner nations in an effort to detect, track, and interdict illegal drugs. All cocaine traffic originates from South America and enters the United States through three main corridors: Central America, Mexico, and the Caribbean islands. JIATF South is the lead agency for counterdrug operations and finds itself being asked to do more with less against drug organizations with billions of dollars in revenues (Insulza, 2012). Due to budget cuts and the decrease in the number of surface assets available to conduct drug interdiction operations, the quantity of illegal drugs intercepted has declined in recent years: from 240 metric tons in 2010 to 152 metric tons in 2012 to 132 metric tons in 2013 (Freedburg, 2013). This thesis examines probability and optimization models for the maritime domain that will help counterdrug operation planners to maximize the returns on their limited resources.

A. BACKGROUND

Since the mid-1970s, the “United States has invested billions of dollars in counterdrug programs,” with the objective of reducing the flow of Latin American-sourced illicit drugs into the United States (Seelke, 2010, p. 1). Past interdiction methods included stricter cross-border control and targeting crop sources in Latin America (Seelke, 2010). Transnational criminal organizations (TCOs) try to exploit land, sea, and air means to smuggle drugs into the United States. Recently, drug smuggling in the southwest region has been dominated by land. Due to increased border security measures in the southwestern United States, counterdrug agencies are finding more evidence of Mexican drug smugglers exploiting maritime routes to get drugs into the country (Seelke, 2010). Drug smugglers from South America are also exploiting maritime means to export illicit drugs to the United States. Drug smugglers from Latin American countries transport cocaine to the United States along maritime routes through the use of container ships, commercial fishing vessels, recreation vessels, and go-fast boats (Seelke, 2010). Traffickers also have increasingly used self-propelled semisubmersibles (SPSSs) to

transport cocaine from South America to Mexico (Office of National Drug Control Policy, 2013b). The use of SPSSs allows traffickers to covertly transport large quantities of drugs, making it more difficult for counterdrug units to detect these shipments (Office of National Drug Control Policy, 2013b). Figure 1 depicts an SPSS; its low profile makes it very difficult to detect.



Figure 1. Self-Propelled Semisubmersible (from *On Patrol Magazine*, 2013).

JIATF South is responsible for the detection, monitoring, and interdiction of vessels and aircraft transporting illicit drugs. Over the last several years, individuals from the Naval Postgraduate School (NPS), Naval Research Laboratory (NRL), Sandia, Space and Naval Warfare System Command (SPAWAR), and the University of Connecticut (UConn) have worked to develop an optimized search planning tool that recommends search plans against drug-trafficking organizations in the maritime domain and maximizes the use of available assets. This thesis focuses on searching and interdicting TCOs in the maritime domain.

South American, Caribbean, and Mexican TCOs are responsible for most large-scale maritime shipments of drugs smuggled directly into the United States (National Drug Intelligence Center, 2011). They typically smuggle cocaine, heroin, and marijuana into the United States through three main corridors: Mexico, Central America, and the Caribbean islands (National Drug Intelligence Center, 2011). Figure 2 illustrates these

three primary corridors used by various TCOs to smuggle cocaine into the United States. Cocaine and marijuana conveyances generally use cargo ships and maritime containers destined for ports in Florida, New Jersey, and New York (National Drug Intelligence Center, 2011). Caribbean traffickers use noncommercial vessels to smuggle cocaine and marijuana into South Florida from the Bahamas and to Puerto Rico from the Dominican Republic (National Drug Intelligence Center, 2011). Mexican TCOs use small noncommercial maritime vessels, commonly referred to as “lanchas” or “pangas,” to transport marijuana to the shores of South Texas and Southern California (National Drug Intelligence Center, 2011). This maritime smuggling method is increasing as land border security strengthens along the United States and Mexican borders (National Drug Intelligence Center, 2011).



Figure 2. Movement of Cocaine into the United States (from Customs Border Protection [CBP], 2014).

To combat increased illicit maritime activity, JIATF-South, along with the United States Coast Guard (USCG), Customs Border Protection (CBP), and federal, state, local, and international partners continue to coordinate standing operations to combat such activity along the country’s maritime borders. These operations include Operation Martillo, along the coastal regions of Central America; Baja Tempestad and Blue Tempest in Southern California; and the Border Presence II and South Texas Campaign in the Gulf

of Mexico (Office of National Drug Control Policy, 2013b). It remains of paramount importance to ensure the effectiveness of CBP and USCG programs to enhance maritime domain awareness and interdict targets transporting narcotics through the maritime domain. The research described in this thesis has the potential to greatly enhance maritime domain awareness by providing planners and search/interdiction assets with better situational awareness of the drug-smuggling activities being conducted within their entire area of responsibility (AOR).

In an effort to assist JIATF South in combating TCOs, we developed and evaluated quantitative tools and methods to increase the effectiveness of counterdrug operations. This thesis builds on the previous work completed by Mooshegian (2013), and seeks to develop simplified and improved methods to determine information about a drug trafficker's course and speed from probabilistic inputs. It also develops a tool prototype that combines several separate modeling components and provides user friendly inputs and outputs. In the remainder of this chapter, we discuss in more detail the key TCOs involved with the maritime drug trade, the U.S. players involved in the maritime domain, maritime law, and the societal cost of illegal drugs in the United States.

B. MARITIME TRANSNATIONAL CRIMINAL ORGANIZATIONS

Colombia, Dominican Republic, Cuba, and Mexico dominate the maritime routes used to traffic illicit drugs directly into the United States. The Colombian transnational criminal organizations (TCOs) are largely remnants of larger cartels, and smuggle most of their drugs into eastern U.S. markets (National Drug Intelligence Center, 2011). Colombian traffickers generate profits by selling cocaine and heroin to Mexican and Caribbean traffickers for distribution in the United States.

Dominican TCOs in the United States have long-standing associations with Colombian TCOs and have increased their prominence as drug traffickers as they continue to develop and expand trafficking connections (National Drug Intelligence Center, 2011). The Dominican TCOs obtain cocaine and heroin from TCOs in Colombia and smuggle the drugs into the United States (National Drug Intelligence Center, 2011). The involvement of Dominican TCOs in domestic drug trafficking is projected to

increase in the near future as expand trafficking connections (National Drug Intelligence Center, 2011). Dominican trafficking organizations operate independently of one another; there is no national or regional leadership (National Drug Intelligence Center, 2011). Consequently, the expansion of these organizations depends on the exploitation of local opportunities (National Drug Intelligence Center, 2011).

Cuban trafficking organizations and criminal groups have expanded their drug-trafficking activities beyond the Florida/Caribbean Region (National Drug Intelligence Center, 2011). This expansion evolved as a reaction to increased law enforcement pressure in South Florida. Cuban traffickers primarily distribute high-potency marijuana in Florida (National Drug Intelligence Center, 2011).

Mexican-based TCOs control much of the production and distribution of illicit drugs destined for the United States (National Drug Intelligence Center, 2011). Smugglers move most of the drugs through the southwest border (National Drug Intelligence Center, 2011). Increased border enforcement forces Mexican TCOs to rely more on alternative smuggling methods such as using noncommercial maritime vessels, commonly referred to as lanchas or pangas (National Drug Intelligence Center, 2011). Figure 3 illustrates video surveillance taken by the USCG of a panga off the coast of Southern California. According to the National Drug Intelligence Center, maritime smuggling is anticipated to increase as land border security continues to be strengthened along the U.S.-Mexican border (National Drug Intelligence Center, 2011). This thesis will focus on TCOs maritime smuggling routes into the United States, originating from South America.



Figure 3. Panga Smuggling Narcotics off the Coast of California (from Military.com, 2014).

C. UNITED STATES COUNTERDRUG ORGANIZATIONS

United States Southern Command (SOUTHCOM) is the lead U.S. agency responsible for directing illicit trafficking detection and monitoring activities (Kelly, 2014). SOUTHCOM has found itself relying on increased contributions from partnering nations in order to achieve mission success, but not without degradation, as force allocation cuts by the services have taken a toll on operational results. Drug seizures dropped from 240 metric tons annually in 2010 to 132 metric tons annually 2013, due to limited assets (Kelly, 2013). SOUTHCOM's lack of assets will continue to constrain the operation's effectiveness. To mitigate asset shortfalls, SOUTHCOM relies heavily on the USCG and CBP, which now provide the bulk of the ships and aircraft available to disrupt drugs bound for the United States. When then-chief of SOUTHCOM, Air Force General Doug Fraser, spoke to the Defense Writers' Group, he estimated that JIATF South was only able to intercept about 25 percent of the drug shipments they knew about and were able to track (Frasier, 2010). This low rate derives primarily from a lack of available assets (Freedberg, 2013). The work done in this thesis directly addresses how to more effectively allocate limited assets in support of SOUTCOM's mission.

1. Joint Interagency Task Force South

Joint Interagency Task Force South (JIATF South) serves as the center for counter-illicit-trafficking operations (Kelly, 2013). JIATF South comprises representatives from the military, government agencies (e.g., Department of Homeland

Security, Federal Bureau of Investigation, Central Intelligence Agency, and Drug Enforcement Agency), and foreign partners. A number of South and Central American countries have also assigned liaison officers to JIATF South in an effort to combat illegal drug trafficking (Office of National Drug Control Policy, 2013). In 2013, international and cooperative interagency efforts coordinated through JIATF South resulted in the seizure of over \$2.6 billion worth of drugs and the detention of 295 suspects (Kelly, 2013). The work of this thesis will directly impact JIATF South in the planning of counterdrug search and interdiction planning.

2. United States Coast Guard

Founded in 1915, the U.S. Coast Guard (USCG) is the primary maritime law enforcement agency of the United States. The Coast Guard can make inspections, seizures, and arrests on waters within the United States' jurisdiction and jurisdiction of partnering nations. The Coast Guard is the lead U.S. agency for maritime drug interdiction and has many assets that are beneficial to counterdrug operations including major cutters, fixed-wing and rotary-wing aircraft, and patrol boats. Major cutters can provide a presence at sea to respond to smugglers far from the coasts. Cutters have the capability to detect, interdict, stop, and board other vessels. The Coast Guard also has fixed-wing surveillance aircraft that can launch from air stations to conduct search patrols. Medium- and short-range response helicopters are another asset that the Coast Guard can deploy from air stations to conduct interdiction operations. The Coast Guard remains one of the most active organizations in assisting JIATF-South in search and interdiction operations.

3. United States Navy

The U.S. Navy has been conducting law-enforcement operations with Coast Guard personnel against drug traffickers for many years (United States Naval War College, 2011). With the recent decline of available naval assets, it is very important that the U.S. Navy continues to work alongside the USCG to create synergies in counterdrug operations. The Navy had to accept a limited role in counterdrug operations over the past few years due to the decommissioning of the Oliver Hazard Perry-class frigates (FFG)

and the slow production of the littoral combat ship (LCS). By the end of 2015, the Navy plans to retire all of its frigates (Larter, 2014). Currently, the Oliver Hazard Perry-class frigates play an oversized part in Navy counterdrug operations and the Navy still has questions on whether the LCS will have the capacity to participate in counterdrug operations (Larter, 2014). The retirement of all the frigates and all of the uncertainty involving LCS will severely limit the Navy's role in counterdrug operations and impact JIATF South in a negative manner. The implementation of decision tools could relieve the stress on JIATF South caused by the decline in the number of available assets and become even more powerful once additional assets become available in the future.

D. LAW OF THE SEA

Understanding the rules set forth in the United Nations Convention on the Law of the Sea (UNCLOS) is vital to counterdrug organizations. From a modeling perspective, the optimization model recommends where assets should go for search and interdiction purposes. From an operational perspective, there are several issues that need to be considered prior to engaging in counterdrug operations. The first is the capacity that search and interdiction assets operate under. When U.S. Navy assets operate in drug-smuggling operations they do so under a law-enforcement capacity and not a military capacity. In order for a U.S. Navy ship to conduct interdiction operations it must fly a Coast Guard flag and operate under a law-enforcement capacity. A second operational issue is having authorization to interdict a vessel of another country. United Nations Convention on the Law of the Sea (UNCLOS) provides the legal basis and guidelines necessary to legally combat illegal drug trafficking (Allen, 1989). UNCLOS streamlines the lengthy diplomatic process required to obtain flag state authority for law-enforcement actions against foreign suspect vessels on the high seas (Allen, 1989). The United States has developed comprehensive agreements to enable maritime interdiction forces to work more effectively and efficiently with other nations. International law permits any nation that has reasonable grounds for believing that a vessel is engaged in illegal drug trafficking to take appropriate action with regard to that vessel (U.S. Department of the Navy, 2007). As a result, maritime counterdrug agreements make territorial boundaries transparent to law enforcement, as they are to smugglers

seeking refuge from interdiction (Allen, 1989). Lastly, an asset's situational awareness of where they are located in relation to all of the legal boundaries is an important operational consideration. According to international law, all international airspace is open to the aircraft of all nations. International waters include contiguous zones, exclusive economic zones, and high seas. Figure 4 illustrates all of the different boundaries associated with the maritime domain. The knowledge and understanding of all the legal constraints is very important because if a search or interdiction asset were to pursue drug smugglers in an area that does not have an established partnership, then the pursuit of one of its flagged vessels could be interpreted as a hostile act against that nation, which would have negative impacts in the international arena. Our planning tool provides operational planners with an advantage by allowing JIATF South planners to know the search plan in advance, which will give them the extra time required to get permission to go into areas that require such authority.




Figure 4. Legal Boundaries of Oceans and Airspace (from United Nations, 1982).

E. IMPACT OF ILLICIT DRUGS ON THE UNITED STATES

Drug trafficking and the use of illicit drugs in the United States has a significant impact on our country. These negative effects include lost productivity, an increase in

crime, an increase in health care costs, and a major strain on our criminal justice system (National Drug Intelligence Center, 2011).

The National Drug Intelligence Center (NDIC) estimates that drug abuse costs the nation more than \$120 billion per year in lost productivity (National Drug Intelligence Center, 2011). A significant factor driving the lost productivity cost is reduced labor participation, which costs society an estimated \$49 billion each year (National Drug Intelligence Center, 2011). While reduced labor participation accounts for missed time at work, it does not account for incarceration. Not only does illicit drug use contribute to reduced productivity, it also has a huge negative impact on our health care system. The NDIC estimates yearly drug-related healthcare costs to be more than \$11 billion (National Drug Intelligence Center, 2011). Additionally, the NDIC estimates the annual cost of drug-related crime at more than \$61 billion. Criminal justice system costs include more than \$56 billion related to illicit drug use. The Arrestee Drug Abuse Monitoring Program (ADAM II), which monitors drug testing among arrestees in 10 cities across the United States, also shows a strong correlation between drug abuse and criminal activity (National Drug Intelligence Center, 2011). The debate about legalization, and whether societal costs will increase or decrease if legalization occurs, is outside the scope of the thesis. Currently, the drugs discussed here are illegal and the government has deemed it a top priority to combat the drug smugglers. This thesis supports that mission by focusing on improving the effectiveness of counterdrug operations.

F. MODELING APPROACH

Over the last several years, a research team has worked on developing a tool to provide situational awareness about the likely location of drug traffickers and recommendations for search plans for counterdrug assets for JIATF South. The team consists of researchers from NPS, UCONN, NRL, and SANDIA National Laboratories. Currently, the JIATF South tool has two main modeling components: a probability model and an optimization model.

The probability model has two main inputs: environmental factors and gathered intelligence. The environmental inputs come from METOC forecasts over a 72-hour horizon. The intelligence input consists of waypoint locations, departure times, velocity, and drug loads. There are also uncertainties associated with the intelligence inputs. For example, the intelligence may point to a vessel leaving from a 100-kilometer (km) stretch of coast, controlled by a particular cartel, during a 72-hr period. As output, it produces the drug smuggler's presence probability as a function of latitude, longitude, and time. We refer to this spatial-temporal presence probability as a heat map. This probability model was developed by NRL (Hansen et al., 2011) and provides a heatmap over a 72-hr window. Figure 5 illustrates what the probability looks like at a specific time for one target.




Figure 5. Snapshot of a Probability Heat Map.

The optimization model developed by Pietz and Royset (2013) produces a search plan for the optimum employment of counterdrug assets. The optimization model directly takes the intelligence information listed above as its input and then outputs an optimal search plan. Thus, there is no connection between the probability model and the optimization model. A connection between the two models may not be necessary for the optimization model to produce operationally effective plans. The probability model, however, accounts for environmental factors, while the optimization model does not. If the weather and environment significantly impact where the drug smugglers can travel,

then connecting the probability model to the optimization model may significantly improve the optimization model's results.

The first phase of this thesis is to develop a rough tool that combines a simplified probability model with the optimization model. The goal is to make the models more user friendly, both for operators and future researchers. The first step is to add an interface for easy user input that only requires information that an operator would have access to. The actual inputs required for the optimization algorithm are somewhat esoteric, and a new user with limited technical background would have a very difficult time using the algorithm. We want to create a buffer between the user and those technical details, so the user only has to enter in standard inputs, such as waypoints and velocities. The outputs to the optimization algorithm are similarly difficult to parse for a non-expert. Thus, we will also create a conversion function that will display the output in a more user-friendly form that novices will understand. This tool, with simplified input and output functionality, will make the models much more accessible. Once completed, we will have a tool that has straightforward input and outputs and can display the results of both the optimization and probability models. Specifically using Google Earth to display the results provides an informative visualization.

The second phase of the thesis builds on the work of Mooshegian (2013). As stated earlier, the optimization algorithm is currently independent of the probability model. If environmental factors play a significant role in the routes taken by smugglers, then this independence could be an issue. Mooshegian (2013) made an initial effort to develop a conversion algorithm that would transform the heat map generated by the probability model to the input for the optimization model. We focus on testing and improving the Mooshegian (2013) conversion algorithm and examine the impact that environmental factors have on a search plan.

G. THESIS STRUCTURE

Chapter II features a literature review of previous work that related to the work in this thesis. Chapter III defines the generation of probability heat maps, search plans, and how we made the computational interface of these models more user-friendly. Chapter IV

discusses the methodology involved in deriving the inputs needed for the optimization model from the probability model. Finally, Chapter V reports the results and Chapter VI concludes.

THIS PAGE INTENTIONALLY LEFT BLANK

II. LITERATURE REVIEW

This thesis uses many components of operations research including search theory, probability, regression analysis, and optimization. In this chapter, we present research relevant to our thesis for each of these areas. We first give a general overview of similar work in search theory in Section A, and then we look at two specific works that thesis builds on. Pietz and Royset (2013) develop the optimization model that generates the search and interdiction plan. Mooshegian (2013) made a first attempt to convert the output to the probability model to the input to the optimization model.

A. SEARCH THEORY

At the core our problem is a search problem: detection and interdiction of drug smugglers. Search theory is one of the oldest areas of operations research, tracing its roots to World War II. The earliest developments in search theory were made by Bernard Koopman and his colleagues in the Anti-Submarine Warfare Operations Research Group of the U.S. Navy during World War II (Nunn, 1981). The original purpose of search theory was to aid the Navy in finding efficient ways to search for enemy submarines and is now relevant to many different applications. There are many military and nonmilitary applications that benefit from search theory, such as looking for enemy submarines, the Coast Guard searching for vessels in distress, prospectors surveying for mineral deposits, forest rangers looking for missing backpackers, or border protection officers searching for drug smugglers (Nunn, 1981). Washburn (2002) provides many mathematical techniques to determine the probability of detecting a target. All search theory-related problems have two elements in common: a desired target and a searcher. There are traditionally two types of costs involved in search problems (Nunn, 1981). The first type consists of costs directly related to the search being conducted (Nunn, 1981). These costs may be measured in dollars, time, manpower expended, or fuel expended. Sometimes these costs are modeled as a constraint. For example, we want to route an aerial search asset to look for a target given the asset only has 10 hours of endurance. A second cost is the cost of not finding the target (Nunn, 1981). This cost may be measured in dollars, in

inconvenience, damage, or lives lost. Sometimes this second cost is framed instead as a reward for finding a target. The two costs need to be balanced in each search situation. In some search problems other costs may appear. For example, there may be a false-positive cost to identifying a nontarget as a target, which would result in collateral damage.

There are many variants and papers related to search theory problems; for example, Slootmaker (2011) applied search theory to combat against piracy. Slootmaker (2011) thesis examines a probability model that estimates the probability of a pirate attack at various locations and times. The work conducted by Slootmaker (2011) is very relevant to the work conducted in this thesis because similar probability models are used to generate for maritime drug smugglers. Johnston (1995) applies search theory to provide aircraft tasking authorities with accurate estimates of detection probabilities for different size search areas, using the surface traffic characteristics and predicted sensor performance for the area of operations. The work conducted by Johnston (1995) is relevant to this thesis because it allows planners to achieve a desired level of surveillance effectiveness for a given-sized search area and considers general traffic flow. While the models described in this thesis do not account for general commercial shipping, it is an area for future work. Pfeiff (2009) applies search theory to a defender attacker optimization model that maximizes the defender's probability of successful detection and classification of SPSSs. Pfeiff's (2009) work is very relevant to this thesis because TCOs utilize SPSSs to smuggle illicit drugs. Pfeiff's thesis takes a more strategic view of the problem, whereas we examine the problem from a tactical view, considering daily or even hourly decisions.

B. SMUGGLER SEARCH PROBLEM

Pietz and Royset (2013) developed an algorithm that generates a search and interdiction plan for the types of problems faced by JIATF South. In this thesis, we consider the search problem introduced by Pietz and Royset (2013), coined the smuggler search problem (SSP). The SSP creates an optimal search plan for searchers routed in an area of interest (AOI) to detect multiple targets moving in a piecewise-linear fashion. This model focuses on routing aerial searchers within an AOI to detect and monitor

targets. The model also positions surface interdictors, but we will not consider that aspect in the thesis. Many optimization models would approach this problem by discretizing time and space (see, e.g., Sidoti et al., 2013). One problem that arises with discrete time and space problems is that the problem size can grow very quickly as we refine the discretization. Instead, Pietz and Royset (2013) develop a continuous time and space model. The decision variables correspond to the order to search the targets, how long to search each target, and when to start searching each target.

The SSP is a specific application to a more general class of problems called the generalized orienteering problem with resource-dependent rewards (GOP-RDR) (Pietz & Royset, 2013). The orienteering problem (OP) originates from the sport game of orienteering and is a combination of the knapsack problem (KP) and the travelling salesperson problem (TSP) (Vansteenwegen, Souffriau, Vanden Berghe, & Van Oudheusden, 2009). In this game, individual competitors start at a specified point and try to visit as many checkpoints as possible, and then return to the original point within a given time frame (Vansteenwegen et al., 2009). Each checkpoint has a certain score and the objective is to maximize the total collected score (Vansteenwegen et al., 2009). The OP's goal is to maximize the total score collected, while the TSP tries to minimize the travel time or distance (Vansteenwegen et al., 2009). The GOP-RDR extends this idea by allowing the underlying graph structure to evolve over time so that distances between vertices are not constant. This applies to situations where searchers are looking for moving targets and the targets represent the nodes in the underlying networks. The generalized orienteering problem has been used to obtain optimal mission plans for military intelligence, surveillance, and reconnaissance (ISR) aircraft. Royset and Reber (2009) consider a more general problem by adding considerations for aircraft take-off times, airspace deconfliction, and distinguishing between search and transit. Many heuristics and exact algorithms for solving OPs have been proposed in the literature. The reward can be thought of as the value from detecting the actual target. The resource-dependent reward relates to limited resources such as time, fuel, and money (cost) that is expended by the searcher while performing search actions. In the SSP problem, this relates to fuel consumption during the search (Pietz & Royset, 2013).

C. PROBABILITY MODELS

Recently researchers at NRL developed piracy models that generate probability maps, or “heat maps,” of pirate locations (Hansen et al., 2011). This model was developed to provide operators with better situational awareness about the likely locations of targets in response to the increasing number of pirate attacks. There had been a significant increase in the total number of pirate attacks in recent years, with 239 attacks in 2006 compared to 439 attacks in 2011 (Hansen et al., 2011). The increase in piracy activities motivated the U.S. Navy to develop a software model that integrated intelligence data, commercial shipping routes, and METOC information to predict regions where pirates may be present and where they may strike next. The model outputs consist of a set of color-coded maps designated the pirate attack risk surface (PARS). These surfaces essentially provide the probability that a pirate will attack at a given time and location (Hansen et al., 2011). Figure 6 illustrates a sample PARS heat map output.




Figure 6. Example of a PARS “Heat Map” Output (from Sidoti et al., 2013).

There are many similarities between the pirate problem and the drug trafficker problem. Both use small vessels that are significantly impacted by winds and sea state,

and both models also deal with multiple maritime targets with uncertainty about their location. Hansen extended this model to the drug trafficker search problem and JIATF South planners currently use this probability model to create search and interdiction plans.

The PARS probability heat map plays a very important role in this thesis since the data within the heat map contains all of the information needed to run the optimization model. Unfortunately, that information needs to be extracted from the heat map. The purpose of this thesis is to convert the heat map output into an input that can be fed into the optimization model developed by Pietz and Royset (2013). The ability to bridge the Hansen et al. (2011) heat map with the optimization algorithm may significantly improve the efficacy of both models.

D. MULTIVARIATE ADAPTIVE REGRESSION SPLINES

Mooshegian's (2013) thesis focused on how to create the target's track from a probability map. This information can then be fed as input into the Pietz and Royset (2013) optimization model to determine optimum employment of JIATF South assets (Mooshegian, 2013). The optimization model assumes that the targets travel in a piecewise linear motion at a known constant speed (Pietz & Royset, 2013). Thus, Mooshegian attempted to derive that information from the probability map. Figure 7 illustrates a sample probability heat map in JIATF South's AOR. There is a waypoint by the Galapagos Islands for one of the targets. We need an algorithm to effectively determine that waypoint.




Figure 7. Sample Counterdrug Probability Heat Map.

Mooshegian's pathfinding model approximates the path taken by the target as a sequence of legs, where the collection of all legs produces the target's entire track from departure location to arrival location. He did this by using a statistical regression technique called multivariate adaptive regressive splines (MARS) (Friedman, 1991). MARS fits a piecewise linear regression model to the probability data by automatically selecting "knots," or points of slope changes. Figure 8 illustrates an example with a clear change in slope and a comparison between an ordinary least squares (OLS) fit and the MARS fit. MARS determines the kinks without user input. This is very important for our situation because we want the algorithm to automatically detect the changes in velocity or course. The MARS method is based on a "divide and conquer" strategy, which partitions the input space into regions, each with its own regression equation. For a full description of MARS, see Friedman (1991) and Hastie et al. (2009).




Figure 8. Illustration of MARS vs. OLS Longitude vs. Time for a Target.

THIS PAGE INTENTIONALLY LEFT BLANK

III. DEVELOPMENT OF THE TOOL

We first developed a rough tool for the probability and optimization models that has user-friendly inputs and outputs. We did this for two reasons. One is to allow operators to use the tool. The other is to make the models and algorithms more accessible to future researchers, so that these individuals can more easily modify and extend the models to other situations. The main probability model formulated by NRL is a complicated simulation model that incorporates METOC factors (Hansen et al., 2011). In our tool, we used a much simpler probability model developed by Mooshegian (2013) that mimics the NRL model at a high level. It would be straightforward to substitute the NRL probability model for the Mooshegian model in the tool. While the Mooshegian probability model is quite useful, prior to the work conducted in this thesis, it was not in a form readily usable by others. The model is implemented in Matlab and requires a strong knowledge of Matlab and perseverance to hunt through the code to determine how to properly define the input parameters for a specific case of interest. The optimization model that generates the search plan is also implemented numerically in a way that makes it difficult for anyone but the original authors to use. The optimization code takes complex, comma separated value (CSV) inputs and uses python and general algebraic modeling system (GAMS) to produce the results. Prior to this thesis, the two models were two separate pieces of code. Not only did we make using each of these models numerically much more user-friendly, but we combined them into one tool. This tool takes simple CSV inputs in and generates both the heatmap and search plans in a form conducive for displaying the results in Google Earth. In this chapter, we describe our effort in developing this tool.

A. GENERATING PROBABILITY HEAT MAPS

The inputs to the probability model include intelligence related to departure location, arrival location, waypoints, departure time, velocity, target type, and the quantity and type of drugs transported. We also have uncertainty associated with many of these quantities (e.g., times, locations, velocities). Finally, we have value between $[0, 1]$

that specifies the probability that the target actually exists. Table 1 lists the intelligence data that plays a crucial role in our analysis.

Intelligence Data	Example Case Information Data
I. Departure Location	(lat,lon) = (11, -74.33)
II. Arrival Location	(lat,lon) =(15.25, -83.24)
III. Waypoint	(lat,lon) =(5,-77)
IV. Velocity	(max, min)=(10,20)
V. Value/Load	2500 kg/Cocaine
VI. Certainty	0.5
VII. Target Type	SPSS
VIII. Departure Time	(Earliest, Latest)=(10,25)

Table 1. Sample Data of the Case Information File.

Given this intelligence, the model generates heat maps to provide the operators with better situational awareness about the likely locations of targets, both in the current period and future periods. A heat map represents a target's path within an AOR. Figure 9 illustrates a heat map representation of a target's path within the JIATF-South AOR. This snapshot of the heat map represents the probability that the target is located in a certain area, at a certain time.



Figure 9. Heat Map Representation of a Target's Path.

To produce the heat map, we must discretize the AOR into cells (e.g., 20 x 20 nm). The heat map is really a three-dimensional data structure: at each time, it specifies the probability the target is in a certain cell, defined by a latitude/longitude point located at the center of the cell. Mooshegian (2013) initially developed the probability model and he implemented the model numerically in Matlab. The Mooshegian (2013) model assumes that given a series of waypoints, the target travels in a straight line between the waypoints. The original code assumes that the specific waypoints are chosen uniformly around some midpoint. We modified the code to include other distributions, such as a triangular distribution. Furthermore, we enhanced the code to include an option that has smugglers take a perturbed path from the shortest path. This serves to mimic the impact of weather at a high level.

An issue with the original code is that it is fairly static. Generating a heat map for a new case would require an individual to dig through the Matlab code to change the input parameters. There is no automation for quickly generating heat maps for multiple cases. Unless one has some technical experience and the time to study the code, it would

be difficult to use this code in practice. Thus, we first made the numerical implementation more user-friendly and general to expand the potential user-base of the model.

The tool requires two input CSVs that are transparent and easy to change. The first input file defines options that specify the boundary of the AOI, the grid cell size, and the time period of interest. In practice, users would rarely modify this file. The second input CSV contains the case information related to each target. Each row of the case information CSV represents a target and each column corresponds to intelligence related to the smuggler. The case information file contains all of the information listed in Table 1, with additional information specifying the uncertainty associated with many of those inputs. The data intelligence on the departure point location (row I of Table 1), arrival location (row II of Table 1), and intermediate waypoints (row III of Table 1) determines a target's track. Each waypoint contains three elements: latitude, longitude, and a length of uncertainty regarding the location of that waypoint. For the velocity (row IV of Table 1) we specify a feasible range. The user also enters the target's value in terms of the expected amount of drugs (row V of Table 1) and specifies the vessel type (row VII of Table 1). Finally, the user can also enter a number between 0 and 1 that represents how certain the intelligence community is that the intelligence is legitimate.

There is still room for improvement in the way that a user enters intelligence data into the model. Currently, the user manually enters in numbers into cells and this can be tedious. For example, to enter in two waypoints into a cell the user would enter something like 2.46;-78.07;80|14.05;-91.38;40. Each waypoint is separated by a pipe (|) and elements of each waypoint are separated by a semicolon (;). While cumbersome, this is a cosmetic issue. Future changes to the tool that would make it more user-friendly include adding drop-down menus or prompts that would ask the user for specific data at the start of running the program. Many improvements have been made the model more user-friendly, but the level of detail that incorporates prompts and drop-down menus will be left for future work.

Once the user has defined the case information for each target, he can run the tool to generate the probability heat maps. Figure 10 illustrates an example with a single

target. The colored area represents the likelihood of a target being in a particular area at different time steps of the planning window.

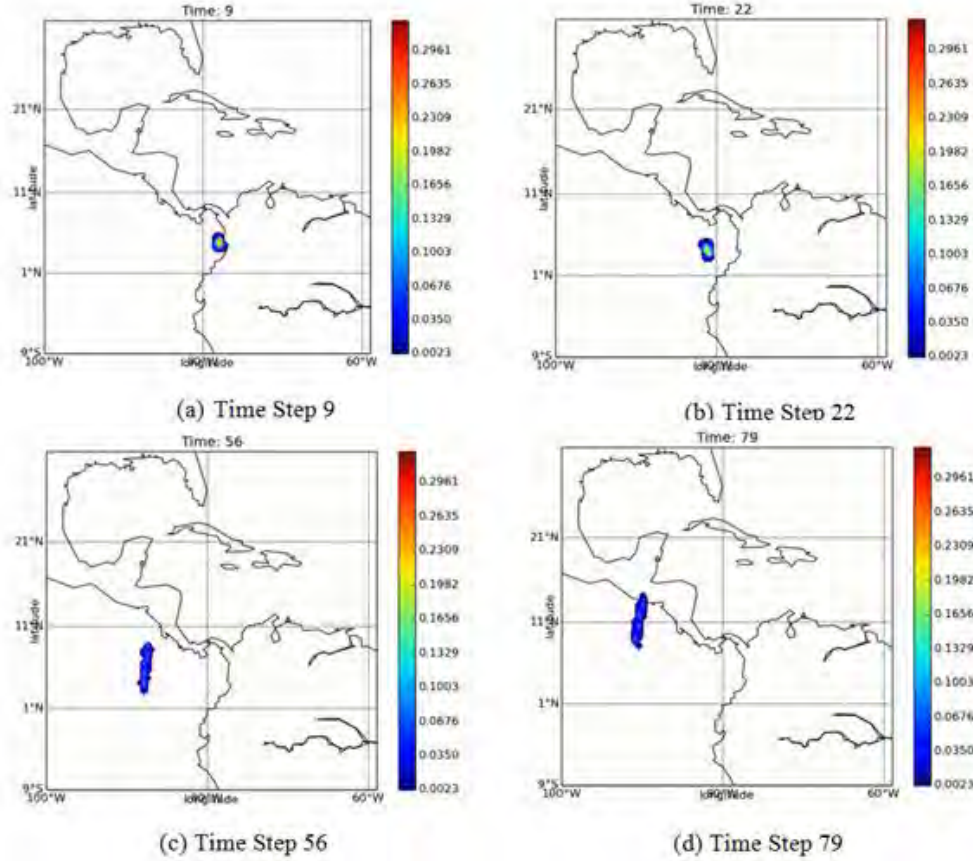


Figure 10. Depiction of the Area of Interest and a Single Target Scenario.

B. OPTIMIZATION MODEL

The probability model described in the previous section should provide great value to operators, as it provides them with a picture for where targets will be over time. However, optimally assigning moving assets to moving targets is a nontrivial task and requires sophisticated optimization machinery. Pietz and Royset (2013) developed the optimization model and there are several flavors of the model that trade-off computation speed vs. optimality gap. Our next step is to incorporate the optimization model into our tool with the probability model so that we can superimpose the search results on the probability heat map.

In practice the operators must coordinate aerial search assets with surface interditors. Once the searcher finds a target, it must notify an interdicator and remain on station until the interdicator arrives. The Pietz and Royset (2013) model does account for this; however, we will ignore interdiction and focus on the allocation of searchers only.

1. Generating Search Plans Using the Optimization Model

The Pietz and Royset (2013) optimization model is a path-constrained optimal search problem in continuous space and time (Pietz & Royset, 2013). The model determines the optimal routing of search assets in the area of operations to maximize the expected drug load seized by the searcher. The optimization model assumes that the targets travel in a piecewise linear motion at a known constant speed (Pietz & Royset, 2013). In real-world scenarios, a target may travel along a track that follows a particular stretch of coastline, or the target may choose to navigate around an island, or maneuver based on the weather conditions or on counter-intelligence. All of these considerations can be approximated with piecewise linear target movement tracks.

The optimization model requires information about the cases of interest. Much of this information is the same as the intelligence used for the probability model described in section A. Specifically, the model needs the information on velocity, waypoints, drug load, vessel type, and expected departure time that appears in Table 1. In addition, the optimization model also requires information on the search assets and their capabilities. Similar to the probability model, it would be difficult for a lay-person to directly run the current numerical implementation of optimization model and interpret its output. While, the model requires standard available information as listed in Table 1, it is entered in a specific format that would confuse most users. The target input file for the optimization algorithm has one segment of one target path on each row. Thus, if a target travels along a path with three piecewise linear segments, the target will appear three separate times in the input file. Our tool creates a buffer between the user and this sort of minutiae. To make the code more user-friendly, we define four separate input CSVs that contain the

relevant information to run the tool. Most users should have this information readily available. Two of these files are essentially fixed and will rarely change. The user will modify the other for each scenario. We describe each of these CSV files below.

2. Options File

The options file contains information relevant to the optimization code's settings that will rarely, if ever, change. The primary option the users might change is the solve method option. This option specifies which algorithm should produce the search plan. On one extreme is a fast heuristic and on the other extreme is an exact branch and bound algorithm. The fast heuristic will not necessarily produce the optimal solution but takes only a few seconds to run. The exact solution can take hours, or longer, to solve, but produces the optimal search plan. There are also several intermediate algorithms between these two extremes that trade-off computation speed versus the optimality gap. In all of the scenarios we examine, the fast heuristic was close to the exact solution. See Pietz and Royset (2013) for a more thorough comparison of the different optimization algorithms. The only other option that a user might change is the objective function. The user may choose to maximize the expected amount of drugs detected or the expected number of vessels detected.

3. Search and Detection Capabilities

The optimization model assumes the searcher performs random search within a search box. We first define T as the random time it takes a searcher to detect a target in the box. The cumulative distribution function for the random variable T is:

$$F_T(t) = P(T \leq t) = 1 - e^{-VWt/A} \quad (1)$$

Equation 1 applies when a targets speed is much less than a search assets speed (V) (Washburn, 2002). The searcher sweep width is W and T is the time of initial detection. We define a sweep-width CSV that lists the sensor capabilities. It specifies the appropriate sweep width for a sensor package against a particular target vessel type. This file is necessary because some drug smuggling vessels are more difficult to detect. For

example a go-fast travels at a high velocity and the majority of the vessel remains above the water making it easy to detect. An SPSS travels at a slow speed and is difficult to detect because the majority of the vessel is submerged with a very low profile in the water. Some sensors do a better job of detecting targets than others. For example, surface search radars onboard surface vessels will perform well in detecting surface vessels. These radars would perform poorly, however, in detecting SPSSs because the waves would obscure the low-profile SPSS from surface search radar. Infrared sensor packages on aerial search platforms perform better in detecting an SPSS due to the heat signature that an SPSS emits. This Sensor Capability CSV is also one that the user should rarely access or modify. The file represents a database that contains every type of sensor available to JIATF South and the detection capabilities against each possible target. The data input into this file contains sensitive information about specific capabilities contained within each asset type and would most likely be updated by a technical expert and not a normal user of the planning process. The user only needs to specify the sensor packages available and the vessel types used by targets. This database would then return the relevant sweep-width. In practice, this database could be made more sophisticated by incorporating aspects such as weather, which can affect the sweep width. We leave that for future work. Table 2 illustrates the sample data located within the sweep-width file. The first column represents various sensor packages (labeled A, B, C . . .) and the associated sweep-width performance against different types of drug smuggling vessels. We made this data up for our research purposes.

	FSV	FV	GF	LPV	MV	PANGA	SNORKELER	SSPS	SV
A	14	12	15	10	20	10	8	5	3
B	10	8	11	6	16	6	4	2	1
C	16	14	17	12	22	12	10	7	5
D	15	13	16	11	21	11	9	6	4
E	12	10	13	8	18	8	6	3	1

Table 2. Sample Sweep-Width Data.

4. JIATF South Search Assets

The next CSV contains search asset information. A user will need to update this for every search scenario. Each row of the CSV file contains information about one asset that would be readily available to any operator. Table 3 illustrates the data contained within the search asset file.

I. Searcher ID	II. Capability	III. Start Location	IV. End Location	V. Speed	VI. Endurance	VII. On Station Speed
1	A	13.93,90.43	24.44,-87.64	325	10	205
2	B	10.08,-72.18	24.85,-83.31	280	8	210
3	A	10.32,-81.67	18.85,-83.92	411	192	328
4	B	29.96,-90.08	16.94,-80.51	117	24	80
5	A	14.12,-78.3	14.12,-78.30	146	4	80

Table 3. Sample Search Asset Information.

The relevant information includes speed of the searcher in traveling to the search area (column V in Table 3) and the on-station search speed (column VII in Table 3). The difference between the speed and on-station speed of the searcher is that the speed parameter is the fastest speed a search asset can travel in order to get from the base to the target and the on-station speed is the speed at which the searcher will conduct its search for the target. We also specify the endurance of the asset (column VI in Table 3), sensor capabilities (e.g., radar package) the search asset has (column II in Table 3). In practice, it might be easier to define a database similar to the search capability file described in the last section. In this way a user could just select the specific asset and the database would automatically fill in the corresponding information. Currently, we do not have this functionality, but it would straightforward to incorporate and is saved for future work.

The final piece of information the user must input is the starting and ending latitude/longitude coordinates of the searcher (columns III and IV in Table 3). Restricting the assets to return to the start location can reduce the effectiveness of the search plan. For example, if a search asset took off from Jacksonville, Florida and searched for a target in the vicinity of Puerto Rico, it does not need to fly all the way back to Jacksonville. The searcher can take off from Jacksonville and plan to land in Puerto Rico

for fuel prior to returning to Jacksonville. Figure 11 Illustrates the assets included in the optimization model broken down by platform type, branch of service, top speed, cruise speed, endurance, and range.







	Type	Service	Top Speed	Cruise Speed	Endurance	Range
	Oliver Hazard Perry Class Frigate	USN	+30 knots	13 knots	7 days	4200 nm
	SH-60 Seahawk	USN	146 knots	60-100knots	4 hours	450nm
	P-3 Orion	USN	411 knots	328 knots	16 hours	2380nm
	Coast Guard Cutter	USCG	19.5 knots	11 knots	14 days	9900 nm
	MH-65 Dolphin	USCG	175 knots	148 knots	4 hours	290 nm
	MQ-9 Maritime Predator	CBP	117 knots	70-90 knots	24 hours	675 nm

Figure 11. All Optimization Model Search Assets.

5. Case Information

The final file contains information about each target. This is the same file used to generate the heat map with the probability model described in Section A and illustrated in Table 1. The operator would need to update this file for every scenario. All of the information on this file will come from intelligence reports and available to the user.

C. OPTIMIZATION OUTPUT

After streamlining the input process to the optimization model, we next turn to the output. The base output of the optimization algorithm is a text file that a user would have difficulty translating into action. The decision variables of the optimization model specify when and where the searcher should arrive to search each target and how long he should remain on-station. We create two output files that will provide insight on the search plan. One file specifies the value of the objective function, which is either the total expected

amount of drugs detected or the total expected number of vessels detected. It also specifies how long each asset should spend searching for each target. This quickly gives the user a summary overview of the search plan illustrated in Table 4.

Total Quantity of Drugs Seized (KG)		3685.1
	Target 1	Target 2
Searcher 1	6.71	0
Searcher 2	6.53	0
Total Search Time	13.24	0

Table 4. All Search Assets Search for Target 1.

The more detailed output CSV specifies the hourly position of each asset during the planning period. Each column of the file represents one hour of the planning period and the rows contain the coordinates of each searcher at the given time. The power of this output file is that it can be used for easy visualization of the search plan. For example, the file can be read into Google Earth where the user can view the proposed plan. By displaying the output on Google Earth, or other similar visualization schemes, users can easily incorporate the proposed search plan into a planning brief. Table 5 illustrates a sample of the detailed output CSV file.

	Time = 5	Time = 6	Time = 7	Time = 8	Time = 9	Time = 10
Searcher 1 Lat	13.93	13.29	10.05	6.9	4.74	4.69
Searcher 1 Lon	-90.43	-89.55	-85.1733	-80.79	-77.93	-78.17

Table 5. Sample Detailed Output CSV File for Searcher's Location over Time.

D. OPTIMIZED SEARCH PLAN OUTPUT

Once all the input files have been properly defined we can run the tool to produce both the heat map and the corresponding search plan. Figure 12 provides snapshots of the output generated from the tool. At hour 5 of the 72-hr planning, a MQ-9 Predator departs

El Salvador for the southern coast of Columbia. At hour 8, the Predator reaches the search area and searches for the drug smugglers until hour 18, when it departs back to El Salvador and arrives back at hour 21.

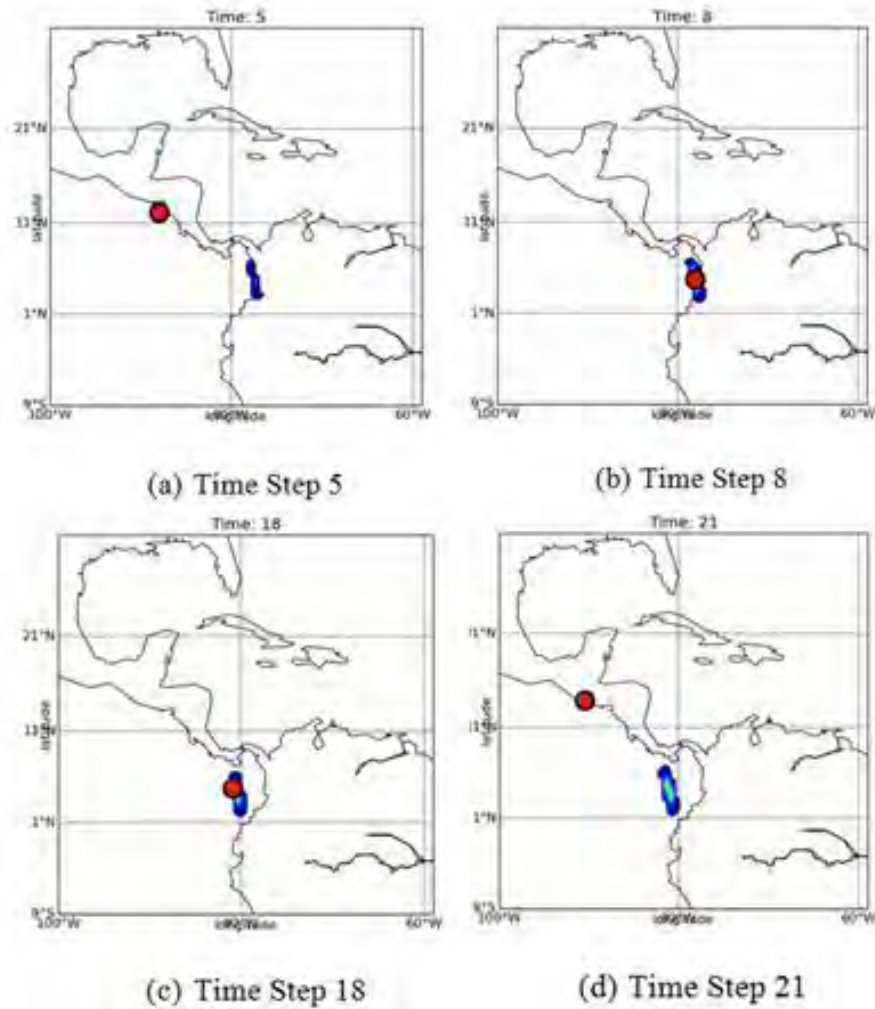


Figure 12. Illustration of the Search Plan.

The blue blob represents the likely location of the target and the red dot represents the location of the searcher. The output of the optimization model gives planners the exact time and exact coordinates for each searcher during the planning period.

E. CONVERSION ALGORITHM

We have now completed a tool that combines the probability and optimization models into a user-friendly tool; however, the two models are still independent. We next focus on generating the inputs needed for the optimization model from probability model heat maps. After we have done this, we can use the tool described in this chapter to easily generate scenarios, test the conversion algorithm, and compare search plans.

THIS PAGE INTENTIONALLY LEFT BLANK

IV. GENERATING OPTIMIZATION INPUT FROM HEAT MAPS

The previous chapter described our effort to combine the probability model and optimization model into a user-friendly tool. The optimization model, however, does not depend on the probability model. In this chapter, we focus on converting the information contained in the heat map into the proper input for the optimization model. As discussed in Chapter III, the optimization model requires four input CSVs; however, three of them will be known by the operator (e.g., assets available and sensor capabilities). The Case Information CSV, however, may not be known by the operator, and we have to generate that CSV from the heat map information. In Table 6, we display one row from the Case Information CSV, and we must extract columns 2-4 from the heat map. Column two contains all waypoints (including starting and ending locations). Each waypoint is divided by a pipe (|) and has three elements associated with it. A semi-colon separates each of the three elements, which are the latitude, longitude, and uncertainty associated with the location, respectively. Thus, we need to determine all waypoints and associated spatial uncertainty associated with each waypoint. The third column specifies the temporal departure time window in terms of the earliest departure time and latest departure time, separated by a semicolon. Finally, column 4 lists the information about the velocity of the target along each leg of the journey. For each leg, we specify the minimum and maximum velocity, separated by a semicolon. A pipe separates information about each leg. If we can extract the information in these three columns from the heatmap, then we can run the optimization model. This chapter describes how we estimate the waypoints, departure time window, and velocity range. We will refer to these estimates derived from the heat map as *code-generated data* to distinguish from the true target information given by the intelligence. To evaluate the effectiveness of code-generated data, we will compare the track it produces to a track generated by the actual underlying intelligence. The testing will be discussed later in Chapter V.

Case	Waypoints	DepTimeRange	VelocityRange	Cert	Value	TargetType
1	5;-77;60 3;-86;100 14.5;-92;100	2;8	10;20 15;30	0.5	2500	SPSS

Table 6. Sample Data of the Case Information File.

The key element of this analysis is a *pathfinding* algorithm that estimates the route taken by the target as a piecewise-linear path. This produces the waypoints and the velocity on each leg. This pathfinding algorithm would work with any heat map, including the NRL and Mooshegian versions. This conversion approach is especially important if the probability model accounts for METOC factors (e.g., the NRL model) because the optimization model cannot account for that directly. The METOC factors taken from weather models can play a significant factor in determining a smuggler's track from the departure location to the arrival location. For example, most maritime smuggling vessels are too small to handle rough seas and, thus, drug smugglers may maneuver to avoid heavy seas or weather conditions may push a smuggler off their theoretical straight-line course. Another key component in our analysis is determining the size of the *area of uncertainty* (AOU) of the target on each segment. This AOU is the size of the “blob” on the heat map. This plays a crucial role in estimating the uncertainty associated with each waypoint. This chapter focuses on improving the pathfinding model developed by Mooshegian (2013) and devising a simpler algorithm to compute the area of uncertainty (AOU). We now describe each of these pieces in more detail.

A. TIMES

We need to specify the departure window of the target: the earliest and latest departure times. This corresponds to column 3 of Table 6. The earliest departure time is the first time there is positive probability on the heat map and the latest departure time is the first time that the sum of the probabilities in the heat map equals 1. In theory, there may never be a time when 100 percent of the probability mass is on the water at one time. This could occur if there is a very large temporal uncertainty regarding the departure time. The optimization model can handle these situations; however, we do not consider these cases here. A separate analysis would need to be performed to develop the proper conversion algorithms for these types of cases.

B. GENERAL PATHFINDING MODEL APPROACH

The optimization model requires that the target travels between waypoints at a constant speed and heading. Therefore, the main purpose of the pathfinding model is to accurately find a piecewise linear path from the probability heat map. To accurately find the target's track, the pathfinding model must detect waypoints where change in heading and/or velocity occur. To do this, we will build on the work that Mooshegian (2013) developed to determine these waypoints. Mooshegian (2013) did not have the benefit of the user-friendly tool we describe in Chapter III. He attempted to work directly with the optimization's original input file and derive those inputs. As described in Chapter III, there are many complex components to those specific input files. By working with the user-friendly input that appears in Table 6 we need to estimate many fewer parameters, and this should produce cleaner, more accurate results. This made the conversion process more straightforward and less cumbersome, thus requiring less work. Our simplified, general, pathfinding model approach utilizes the regression technique called MARS that Mooshegian (2013) utilized in his work. To begin using MARS, we first extract a dependent variable and an independent variable from the heat map data. We describe this process in the next subsection.

1. Expected Track

We first compute the expected track of the target. This is the weighted center of the probability blob for each time step. Figure 13 illustrates a snap shot of a probability heat map at one time step. Each time step is one hour during the planning period. All of the colored areas represent the locations with a positive probability of containing the smuggler at the time of interest. Once the colored region is entirely on the water, the sum of the colored region is equal to 1. The bluer colors correspond to small probabilities and the redder probabilities to larger probabilities.




Figure 13. Snapshot of Probability Heat Map at One Time Step.

For each time period we can compute the expected location of the target, based on the probabilities in the heat map. The AOR is discretized into grid cells and the heat map specifies the probability that the target is at each cell, at each time period. Thus, computing the expected location of the target is straightforward calculation. Figure 14 illustrates a target's path through several snapshots in time.

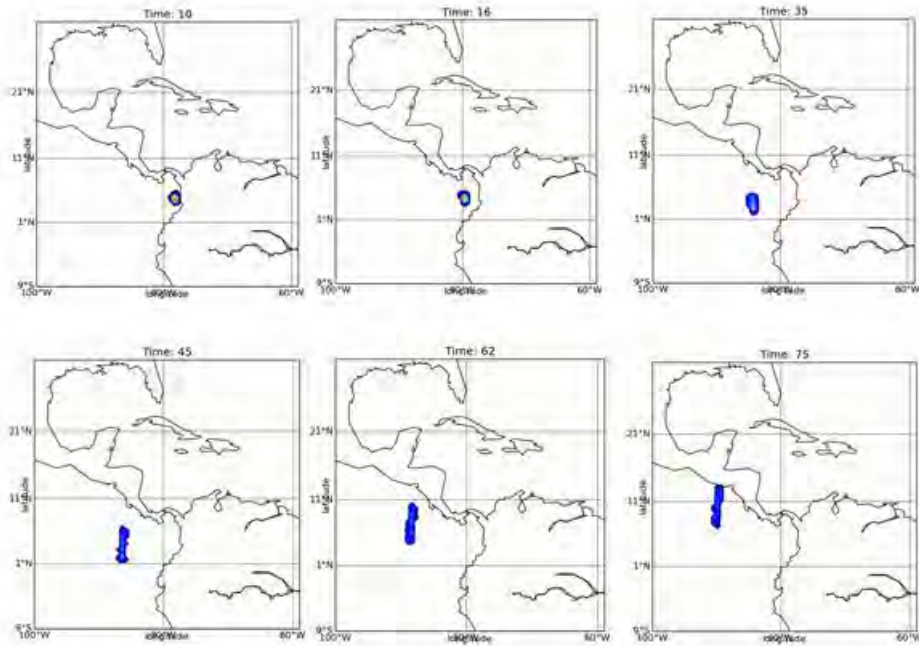


Figure 14. A Target's Track through Time.

Figure 15 plots the expected track for the same target that appears in Figure 16.




Figure 15. Expected Track of Longitude vs. Latitude.

The spacing between successive points captures the velocity. When the spacing between points is close together it indicates that the target is covering less ground at each time step and is moving slower than when the spacing between points is far apart. While difficult to tell, the velocity decreases after longitude -86; the spacing between points decreases. By inspection, a waypoint exists at $3^{\circ} 0' 0''$ North, $86^{\circ} 0' 0''$ West by the sudden change in slope. We will utilize MARS to automate the process of determining the waypoints and creating piecewise linear paths anywhere that a change in speed or direction is detected. It is important to stress that we do want to just capture changes in courses, but also changes in velocity.

2. MARS Function

The MARS model is a generalized regression model (Hastie, 2009). For our purposes, we will use MARS to fit a piecewise linear function to data. It will approximate the path taken by the target as a sequence of piecewise linear paths. The points where a target changes speed or heading changes correspond to a new segment. MARS

automatically fits a piecewise linear regression model to the probability data by selecting “knots,” or points of change, in course or speed. For a full description of MARS, see Friedman (1991) and Hastie et al. (2009).

MARS analyzes the expected track data of the target’s path. The expected track data contains the weighted center of the blob for each time period. Mooshegian (2013) presents the basic MARS model in his thesis and Hastie (2009) gives a much more in-depth treatment. For our purposes, we will illustrate the MARS functionality with a figure. Figure 16 illustrates fitting the expected track from Figure 15, using both simple linear regression and MARS. Clearly a simple linear regression model is inappropriate for this data, but MARS fits the track very well. In many cases, MARS performs excellently in estimating the path of the target without a user needing to specify waypoints prior to running. We aim to systematically test this approach to determine when and how MARS can break down in estimating the path. We discuss the results of this analysis in Chapter V.




Figure 16. Illustrates the MARS Model Fit (Red) vs. the Linear Regression Model Fit (Blue).

Applying MARS to a latitude versus longitude relationship (e.g., figure 16) should determine changes in heading; however, it will not determine points when the velocity changes. Thus, to capture changes in both heading and velocity using MARS we use time as our independent variable and latitude as the dependent, and then do a separate analysis with longitude. Therefore, we combine both “longitude” waypoints and “latitude” waypoints into our final set of waypoints. Mooshegian’s (2013) pathfinding model produced two linear regression equations of the target’s location, one for longitude and one for latitude, which we utilize in our pathfinding model. Figure 17 illustrates how we capture changes in velocity by applying MARS regression techniques for both longitude versus time and latitude versus time. Note that we find a waypoint occurs at time 43 for both longitude and latitude.




Figure 17. Capturing Changes in Heading or Velocity Using MARS.

3. Area of Uncertainty

After generating the waypoints using MARS, we next turn to estimating the area of uncertainty (AOU) surrounding the target at each time period. The optimization model assumes that the velocity and heading are constant on each segment of the target’s path. It also assumes, however, that the AOU is constant on each leg as well. The model makes

this assumption because the AOU represents the size of the search area; the larger the search area, the smaller the probability of detecting the target (see Equation 1). Thus, we must check if the AOU changes significantly on the journey and, if so, appropriately divide constant heading/velocity legs into smaller sublegs with differing AOUs. MARS determines points where a target changes direction or speed; however, since we only consider the weighted centers of the probability blob, we cannot use MARS to determine whether the AOU (i.e., size of the blob) changes between the waypoints. The AOU results from uncertainty in both the departure time and the departure location. Mooshegian (2013) did suggest several ways to estimate the AOU and break up segments into smaller subsegments of constant AOU. While these methods are mathematically sophisticated, they are arguably overkill and did not perform very well in limited testing (Mooshegian, 2013). We start fresh and simplify the approach to estimating the AOU.

We first compute the AOU at each time period. We kept it simple and summed the number of cells with positive probability at each time point. We then multiplied that number by the area of one grid cell to get the AOU for each time. We next need to specify the AOU for each leg of the journey and, if necessary, divide legs into smaller sublegs. The AOU will unlikely remain constant from one time period to the next; at a minimum, there will be small fluctuations. We would like to divide a leg into small pieces only if the AOU changes significantly.

Let us proceed to explain the steps with the help of an example. Assume that the target leaves at time 10 and arrives at time 40, and we determine via MARS that it changes course and/or speed at time 30. We need to define the AOU for the leg on the time interval (10, 30) and a separate AOU for the leg on (30, 40). We may need to further divide those two legs into smaller sublegs if the AOU changes significantly on a leg. We have the AOU at each time unit (e.g., hour) during the time interval of the leg. We first compute the average AOU for the leg. Next, we check if the AOU ever deviates significantly on the leg from that average. We initially set the deviation to 10 percent of the average. If the AOU during (10, 30) remains within 10 percent of the average for the entire time, then we add no new waypoints and define the AOU for (10, 30) as the maximum AOU during that time. We choose the maximum to be conservative, but we

could choose the mean or median AOU on the leg. If the AOU does deviate by more than 10 percent, we divide the leg into sublegs by adding additional waypoints. We want to avoid creating a new waypoint too close in time to an already existing waypoint. Thus, we choose to divide a leg into smaller sublegs based on the length of the time interval of the existing leg. If the time interval is less than 12 hours, we add no new waypoints even if the AOU changes significantly on that interval. If the interval is between 12 hours and 30 hours, we add one waypoint in the middle of the time interval. If the leg is between 30 and 60 hours, we add two equally-spaced (in time) waypoints, and if the leg is greater than 60 hours, we add four equally-spaced points. For each leg or subleg, we always define the leg's AOU as the maximum AOU over the time interval. We could spend a lot of time and energy formulating better methods to divide up legs into sublegs; however, this relatively simple heuristic appears to perform reasonably well.

The Case Information CSV does not take the AOU directly as input. It takes the spatial width of uncertainty in one-dimension at each waypoint. This can be viewed as one side of a rectangle encompassing the AOU. The other side (the length) represents the temporal uncertainty. In order to find this spatial width, we first compute the length of the rectangle. We view the width as capturing spatial uncertainty; i.e., the smuggler leaves somewhere along a 100-km stretch of coast. The length of the rectangle comes from the temporal uncertainty: the smuggler leaves during a 10-hr window. Thus, we can compute the length by multiplying the time uncertainty (latest departure time – earliest departure time) by the velocity of the target. For example, if a suspected go-fast vessel had an earliest departure time of 10 and a latest departure time of 20, and travels at a maximum speed of 45 knots, the resultant length would be 450 nm. Once the length of the AOU is determined, we compute $\text{Width} = \text{AOU}/L$. Figure 18 illustrates an example. The width corresponds to the third value for each waypoint in column 2 of Table 6.




Figure 18. Calculating Area of Uncertainty for a Single Time Step.

4. Determining Waypoint Location

As the algorithm currently stands, it references the location of waypoints by time. Once we have defined all the waypoints based on velocity, heading, and AOU, we need to determine the location of the waypoints. We compute the location of the waypoints by taking the expected location of the target at the time the waypoint occurs. Once we determine all of the waypoint locations for the target's track using our improved pathfinding model, we can display the output on Google Earth. The use of Google Earth to display our pathfinding model outputs enhances our ability to perform analysis by providing a platform to display and examine the results. Figure 19 illustrates our pathfinding model results. The red line represents the target's track, generated from the intelligence used to generate the heat map generator. The blue waypoints are derived by analyzing the resulting heat map and applying our pathfinding model. Our pathfinding model performs well in this example in determining the waypoints along the target's actual track.




Figure 19. Sample Pathfinding Model Output.

Now that we can generate a target's track using our pathfinding model and produce the inputs to the optimization, the final phase of this thesis is to determine how well our improved pathfinding algorithm performs by comparing it to the actual track data used to generate the probability heat maps. Perhaps, more importantly, we also want to examine how the search plans differ between using the actual intelligence and using the code-generated data from the pathfinding algorithm.

THIS PAGE INTENTIONALLY LEFT BLANK

V. RESULTS

The final phase of this thesis consists of two parts. The first is to determine how well our improved pathfinding algorithm performs by comparing the estimated code-generated target track to the actual track data from the intelligence. The second part consists of determining how well our code-generated search plan performs in comparison to the search plan generated using the actual track data. We will consider three measures of performance (MOPs) to evaluate the search plans. The first MOP examines the distance between the search asset and the expected target location for each time step when the search asset actively searches for the target. We refer to this active search time as *time-on-station* (TOS). The second MOP considers the actual TOS for various assets against various targets. The third and final MOP is the expected amount of drugs detected during the search. We utilize Google Earth to visualize the comparisons.

A. CASE GENERATION

To evaluate how well our pathfinding algorithm performs, a total of eight case files are created and built in Google Earth. The eight cases are specifically built to test different aspects of our pathfinding algorithm. These aspects include determining waypoints, calculating the AOU, and adding additional waypoints based on changing AOU. The eight scenarios appear in Figure 20 and define the behavior of the targets. A red line represents a target's actual track based on intelligence. We also consider different combinations of search assets, such as the P-8 Poseidon, SH-60 Seahawk, and MQ-9 Predator, against a variety of targets, such as go-fast boats and SPSSs. We will discuss several of these scenarios in more detail in the next section.



Figure 20. The Actual Tracks of Eight Cases of Interest.

B. SINGLE TARGET RESULTS

In this section, we examine single-target scenarios and primarily focus on situations where our algorithms produce imperfect results. When the target travels at a constant heading, at a constant speed, our algorithms generally perform very well. Thus, we want to examine less “vanilla” scenarios to determine when our algorithms breakdown. These situations will provide material for future work to improve the algorithm. To avoid repetition, we only present detailed results for four of the eight cases mentioned in Section B.

1. Case 1: Straight Path Route with Changing Velocity

The first case file analyzes a go-fast boat departing from Santa Rosa, Colombia with 1,000 kilograms (kg) of cocaine onboard and arriving at Akumal, Mexico. This target travels on a single leg, at a constant heading; however, it reduces its speed approximately halfway through the journey. This case is designed to test the algorithm’s ability to detect changes in speed. Figure 21 illustrates a snapshot of the AOU at a single

point in time. The red dot represents an MQ-9 predator drone with a starting location of San Salvador, El Salvador. The drone has a maximum endurance of 24 hrs.



Figure 21. Snapshot of Area of Uncertainty for Case 1.

We generate the heat map shown in Figure 21 using the actual intelligence data. We then use our pathfinding algorithm on the heat map output to derive the code-generated track. Figure 22 illustrates both the actual track waypoints (red bubble icons and line) and the code-generated track (blue bubble icons). Figure 22 also displays the TOS portion of the search plans. The red plane icons represent the searcher locations during the actual track search plan and the blue plane icons represent the searcher location during the code-generated search plan. We first discuss the performance of the pathfinding algorithm and then compare the two search plans.




Figure 22. Search Plan Results for Case 1.

By inspection of Figure 22, we see that the pathfinding algorithm performs very well in estimating the track of the target. Figure 23 illustrates the MARS fit to the data. The slope change at hour 27 captures the target slowing down; the fit is essentially perfect. Furthermore, the pathfinding algorithm determines that a significant change in the AOU's size during the target's second leg defined by the two red waypoints closest to Mexico. The pathfinding algorithm adds a new waypoint at time step 40 because the AOU changes by a significant amount on the second leg. This new waypoint is the blue center waypoint of the second leg in Figure 22. This occurs because the AOU compresses as the blob transitions from the first leg (with faster speed) to the second leg (with slower speed). As the blob transitions, the part of the blob on the second leg moves slower than the part of the blob on the first leg, which leads to a compression effect that does not finish until later in the second leg. This leads the pathfinding algorithm to add an additional waypoint. Overall, the pathfinding algorithm performs well in determining the target's course and changes in velocity and AOU.




Figure 23. Case 1 MARS Model Fit.

We next turn to the comparison of the search plans. We see from Figure 22 that the plans are very similar, with the code-generated plan starting about four hours earlier. While difficult to see, the red plane icon is directly over the blue plane icon for much of the search. The searcher wants to search at the end of the route because that is closest to the searcher's base. For the code-generated track, this portion also has a smaller AOU. The reason that the two plans do not match up exactly highlights a shortcoming with our approach. Note that the estimate of the arrival location (western-most blue balloon) is slightly off from the actual arrival location. We currently estimate the final waypoint using the last time all the probability mass is on the water. We do this because the optimization disallows any search after that time. Future work should correct for this discrepancy and estimate a better arrival location. The total TOS for the actual plan is 13.04 hours and the TOS for the code-generated plan is 12.84 hours. Both plans effectively track the target during the TOS: the mean difference between the searcher and expected target location is 2.2 during the actual search and 5.6 during the code-generated search. Finally, the code-generated search detects 744 kg of drugs, which is 0.89 of the expected amount detected from the actual search plan (835 kg). Technically, the actual track search plan outperforms the code-generated search plan according to our MOPs. This is not surprising, as the quality of information contained in the actual intelligence

goes through degradation as it flows through the probability model to the heatmap, and then into the pathfinding algorithm. In practice, however, our pathfinding algorithm generates an effective search plan.

2. Case 2: Multiple Leg Route

Case 1 represents a situation where our pathfinding algorithm performs well. Many examples and cases mimic this high-quality performance; however, for the next three subsections, we will discuss situations where the results are not as clean. Finding situations where the algorithm underperforms will help lead to future improvements. The second case file analyzes a go-fast boat departing from Bajo Baudó, Colombia with 1,000 kg of cocaine onboard and arriving to Ocho, Guatemala at a velocity of 15 knots. The search asset assigned to this search is a P-8 Poseidon based out of San Salvador, which has an endurance of 16 hours. Figure 24 illustrates the AOU at two times. At the earlier time, the AOU is small and compact, but at the later time, the AOU has grown significantly. All else being equal, planners would prefer to search for this target early in its journey to take advantage of the small AOU. The AOU increases because spatial uncertainty at later waypoints increases. Figure 25 illustrates the actual track, the code-generated track, and the corresponding search plans. This target heads west from Colombia and then turns significantly to the north between the Galapagos Islands and Costa Rica. This case tests the pathfinding algorithm's ability to determine changes in heading and to detect changes in AOU.

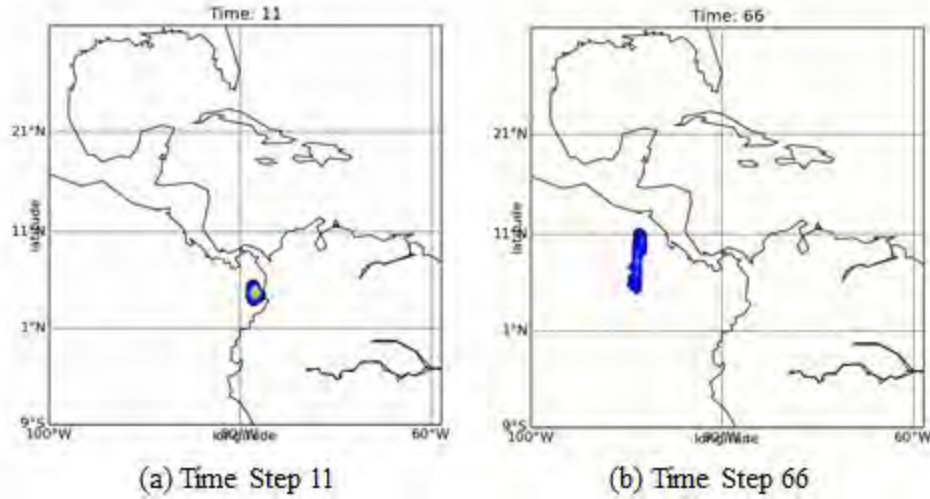


Figure 24. Snapshots of the Heat Map for Case 2.

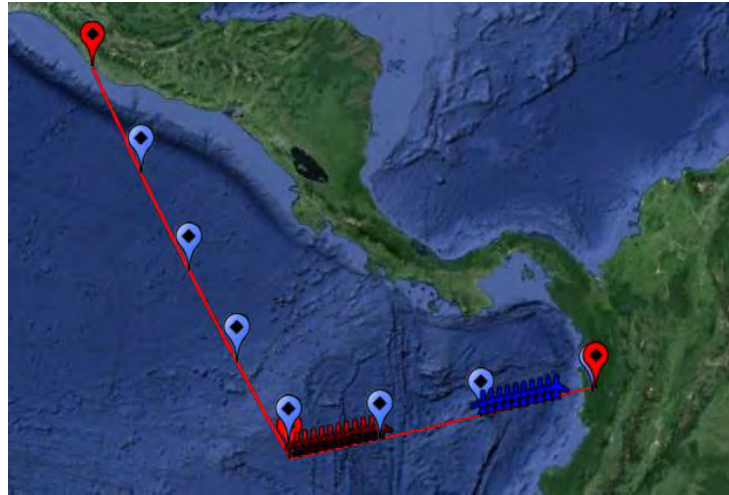


Figure 25. Search Plan Results for the Multiple Leg Case.

The pathfinding algorithm determines the change in heading and also adds additional intermediate waypoints, based on changes in the AOU. While the code-generated track in Figure 25 appears to be a very good fit, a closer inspection reveals a common issue with our algorithm. A “corner cutting” effect occurs near the waypoint where the heading changes. The estimated waypoint is 27 nm away from the true waypoint. We can see this effect more clearly in Figure 26, which presents the MARS fit. MARS fits the longitude values well, but it slightly off with the latitude.




Figure 26. Single Target, Multiple Leg MARS Model Fit.

This corner-cutting effect is a result of the blob “turning the corner” on the heat map. Part of the blob is on one leg and part of the blob is on the other. Figure 27 illustrates this. The larger the blob, the more significant this corner-cutting effect can be.



Figure 27. Snapshot of the Corner-Cutting Effect.

Even though the corner-cutting effect causes the estimated track to deviate from the true track, the code-generated search plan performs very well. During the TOS, on average, the searcher is 2.5 nm from the expected location of the target. The actual search plan is slightly higher, at an average of 5 nm. The code-generated case begins its search

at an early time steps, when the AOU is tight and compact, and the actual search plan begins its search at a later time, when the AOU is wider and is less compact. The actual search plan does not account for changes in AOU as well as the code-generated plan. It naively assumes that the AOU is the same along the entire route before the change in heading. By accounting more precisely for this changing AOU, the code-generated plan has the advantage over the actual plan. The code-generated plan's TOS is 10.5 hours, which is slightly less than the 11.5 hours TOS from the actual search plan. The code-generated plan, however, actually produced a higher objective function than the actual search plan: 745 kg versus 698 kg. The code-generated search plan is far from the change in heading. If the search plan is closer to this waypoint, then the corner-cutting effect may have produced a much less effective search plan.

3. Case 3: Zig-Zag Route

Our pathfinding algorithm performs reasonably well in estimating the route of the targets in cases 1 and 2. Correcting the corner-cutting effect is a future research topic to improve the algorithm; however, in many cases, the effect will have a minimal impact on the resulting search plan. In the next two cases, we will examine situations where MARS breaks down and the estimated track fits the actual track poorly. In case 3, the target zig-zags through the AOR, departing from Puerto, Colombia with 5,000 kg of cocaine and arrives in Usibila, Honduras. The search asset is a P-8 Poseidon and is based in San Salvador, El Salvador, and has an endurance of 16 hours. Figure 28 illustrates a snapshot of the heatmap.




Figure 28. Snapshot of the Area of Uncertainty for the Zig-Zag Case.

Figure 29 illustrates the tracks and search plan results.




Figure 29. Search Plan Results for the Zig-Zag Case.

The code-generated track does not mimic the zig-zagging and just cuts through the AOR. Figure 30 shows the MARS fit by the latitude and longitude. The model performs well for the longitude as the target moves in a general western direction; however, the latitude track estimated by MARS does not capture north/south oscillations well. Furthermore, when we combine the latitude waypoints and longitude waypoints, we

get three waypoints all very close to each other. There are two waypoints on the first leg in Figure 30 near the blue balloon icon; unfortunately, the waypoint icons are covered by the plan icons. Having multiple waypoints in the same vicinity is undesirable. An area for future research would be to “deconflict” these close waypoints and only choose one point.




Figure 30. MARS Model Fit for Zig-Zag Case.

Even though the code-generated track does not fit well, the search plan surprisingly does well. As illustrated in Figure 29, the code-generated search plan closely mimics the actual search plan. On average, the code-generated search plan is 14.5 nm away from the target during the TOS. This is a reasonable amount higher than the 2.5 nm average distance for the actual search; however, the code-generated plan is still effective. For example, an aerial search asset flying at an altitude of only 300 feet can see a radius of 20 nm. The two search plans produce nearly identical TOSs: 13.5 hours for the code-generated plan and 13.2 for the actual plan. Finally, the code-generated search plan detects 4,797 kg, which is 96 percent of the amount detected in the actual search (4,998 kg). While, in this case, the search plan performs very well in spite of the poor track fit, slight variations of these types of cases would produce a very poor search plan. The next section describes such a case.

4. Quick Change of Heading

The next scenario is a case where a go-fast smuggling vessel travels at a high rate of speed for a very short duration and changes course. The go-fast boat departs from La Rada, Colombia with 5000 kg of cocaine onboard and arrives in Berta, Costa Rica. The search asset assigned to the search is a P-8 Poseidon based in San Salvador, El Salvador, with an endurance of 16 hours. Figure 31 illustrates a snapshot of the AOU for the fast target, short duration case.




Figure 31. Snapshot of the AOU for Case 4.

Figure 32 illustrates the search plans.

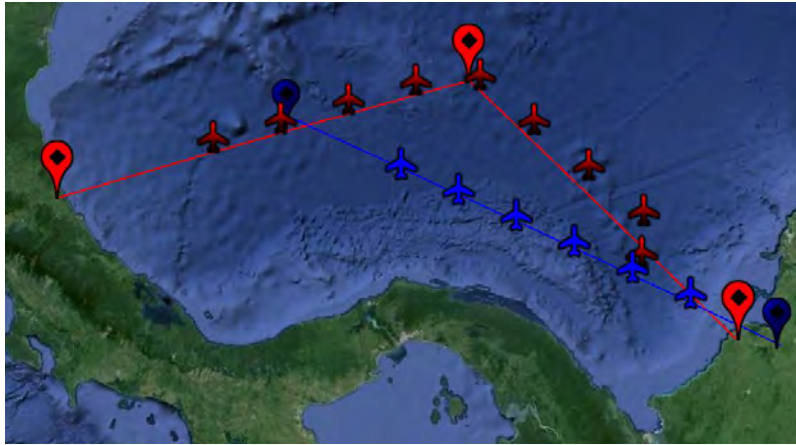


Figure 32. Search Plan Results for Case 4.

We see that the pathfinding algorithm in this case performs terribly and does not find the obvious waypoint. The MARS plots in Figure 33 illustrates that even though there is a clear, nonlinear form to the latitude data, and a piecewise-linear function would fit that data well, MARS returns the OLS line. Comparing this case with the previous one, the MARS algorithm can miss waypoints if the target only remains on a constant heading for a short period of time. In this case, MARS only has eight data points to estimate the track. It is possible that using linear interpolation to generate more points would increase the effectiveness of MARS. This is an area for further research, as the standard MARS algorithm falls short in these cases.




Figure 33. MARS Output for Case 4.

Unlike in case 3, the code-generated search plan also performs poorly. Figure 32 illustrates the searcher is not close to the target under the code-generated search plan. The average distance between the searcher and target is 52 nm. Since the searcher is not close to the target, the other two MOPs do not provide any additional insight.

C. MULTIPLE TARGET RESULTS

In Section B, we consider only one target in the AOR. In reality, there will be many targets. In this section, we examine a multiple target scenario. The multiple target case analyzes a go-fast boat and an SPSS. The go-fast departs from Cartagena, Colombia with 2,000 kg of cocaine and arrives at Tasbarraya, Honduras, while making a heading change once during the journey. The SPSS departs from Cristo Rey, Colombia with 5,000 kg of cocaine and travels at a constant heading to Puerto Cabezas, Nicaragua. Figure 34 presents the AOU and Figure 35 illustrates the tracks and search plans. There are two search assets assigned to handle this case. The first is a P-8 Poseidon (red x) based in San Salvador and the second is a SH-60 Seahawk helicopter (red dot) embarked onboard a frigate underway in the Caribbean Sea. The helicopter icons represent the SH-60.




Figure 34. Snapshot of the Targets Area of Uncertainty for Case 5.

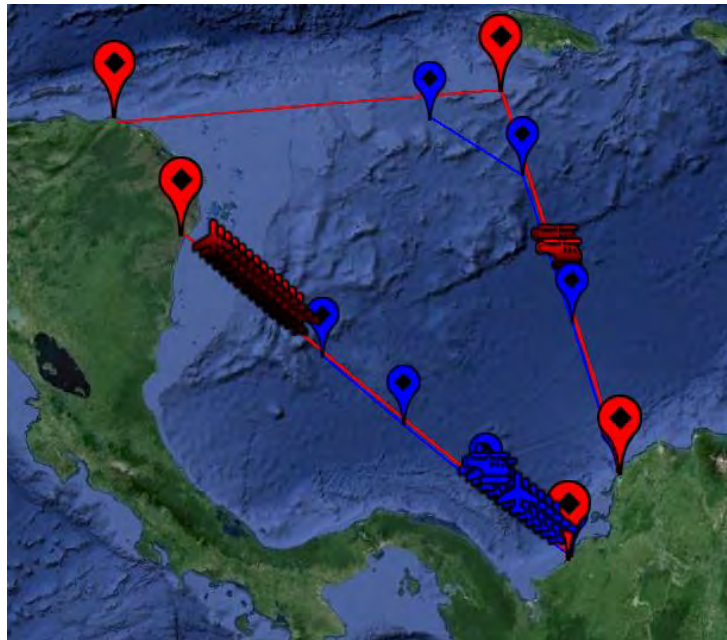


Figure 35. Search Plan Results for Multitarget Scenario.

We see the corner-cutting effect in the go-fast boat target the changes heading. This is the same phenomenon that we discussed earlier for case 2. The code-generated track for the constant-heading SPSS is essentially a perfect fit. Both targets have an uncertain velocity (e.g., the SPSS travels at velocity uniformly chosen between 5 and 15

nm/hr). This causes the blob to stretch over time, as the “front” of the blob represents the fastest possible velocity and the “back” of the blob represents the slowest possible velocity. Thus, the AOU increases over time and the pathfinding algorithm adds additional waypoints (as illustrated in Figure 35) to account for this.

The actual search plan sends a SH-60 asset to search the go-fast target for 1.59 hrs and a P-8 Poseidon asset against an SPSS for 12.89 hrs. The code-generated search plan sends a SH-60 asset to search the SPSS target for 1.26 hrs and also sends the P-8 Poseidon asset against the SPSS for 12.16 hrs; neither asset searches for the go-fast. The code-generated search plan searches earlier in the SPSS’s route because there is a smaller AOU and it gets much larger later in the route. The actual search plan is not accounting for this changing AOU because it assumes a constant velocity, rather than a random velocity. Hence, the actual plan searches the SPSS much later in the route because the actual plan mistakenly thinks the AOU is constant along the entire route. This illustrates an important situation, where the pathfinding algorithm adds value to the optimization. The optimization requires a constant velocity and, if that is not the case, then the actual search plan may produce poor results. Both search plans are, on average, less than 5 nm from the targets during the search. The code-generated search plan detects much less than the actual search plan (1,434 kg versus 3,700 kg) because the code-generated plan only searches for the SPSS. This number is somewhat misleading, however, because the actual search plan assumes that the AOU is much smaller than it actually is.

Most applications of the optimization algorithm only consider land-based assets, such as P-8s. The current version of the optimization model allows a search asset to conduct only a single search during an entire 72-hr planning period. There are more advanced versions of the optimization model that do consider multicycle planning. A future area of research would be to incorporate that model into our tool. The multicycle planning for helicopter search assets would be particularly useful. Helicopters have the capability to refuel while at sea and can conduct continuous flight operations at sea by swapping out flight crews. Multicycle optimization models would allow helicopters to

conduct multiple searches during a 72-hr planning period and would tremendously improve the overall performance due to the helicopters flight endurance of four hours.

D. COMPARISON ACROSS CASES

We conclude this section by comparing our measures of effectiveness across multiple cases. We have eight cases illustrated in Figure 20. We discuss four of the cases in detail in Section B. The time-on-station times were collected for all eight of the case files developed. We examine the difference between actual search plan TOS and code-generated search plan TOS. The actual search plan outperformed the code-generated search plan by having a slightly higher TOS in all eight cases analyzed. On average, the actual search plan spent an hour more on station than the corresponding code-generated plan. Figure 36 illustrates the difference between the actual search plan and code-generated search plan in boxplot form. The black line in the boxplot is the median of the data. The box represents the middle 50 percent of the data and the whiskers represent an estimate of the valid range of the data.




Figure 36. MOP-2 Difference Between Time-On-Station Results.

Our final MOP is the expected amount of drugs detected during the search. To compare the actual search plan to the code-generated search plan we define the code-generated search plan as a percent of the actual search plan. On average, the code-generated plan detects 75 percent of the drugs detected by the actual search plan. Figure 37 illustrates the results for the eight cases. The results in this section show that while there are issues with the pathfinding algorithm, the current version will usually produce effective search plans. Hopefully, future research can address some of these issues to produce a better algorithm.




Figure 37. MOP-3 Percentage of Drugs Seized as a Result of the Search Plan.

VI. SUMMARY AND FUTURE WORK

A. SUMMARY

Over the last several years, individuals from the NPS, NRL, SANDIA, SPAWAR, and UCONN have worked to develop an optimized search planning tool that will recommend search plans against drug-trafficking organizations in the maritime domain and maximize the use of available assets. This thesis supports that effort by examining probability and optimization models that will help counterdrug operators maximize the returns on their limited resources. The first phase of this thesis develops a tool that combines a probability model with the optimization model. Prior to this thesis, the probability and optimization models were implemented in two separate pieces of code that were difficult for any nonexpert to use. The goal is to make the models more user-friendly both for operators and future researchers. While both the probability and optimization models now exist together in a user-friendly tool, the two models are independent. In the second phase of the thesis, we focus on generating the inputs needed for the optimization model from probability model heat maps. This analysis builds on the work of Mooshegian (2013). We use a statistical regression technique called multivariate adaptive regressive splines (MARS) to determine waypoints where the target changes heading or speed.

Next, we evaluate how well our pathfinding algorithm performs by examining both the resulting tracks and search plan. The performance of the algorithm in producing accurate tracks depends on accurate estimation of waypoints. In most situations, the estimated pathfinding track matched well with the actual track. In some situations, the pathfinding algorithm produces a “corner-cutting” effect when the target changes heading. The algorithm correctly determines the waypoint, but the exact location of the waypoint is slightly off. The algorithm can also perform poorly when the target changes heading significantly after a short period of time. We do not have enough information in these cases to accurately determine that change.

We also tested the search plan generated by the pathfinding algorithm using three measures of performance: total distance between searcher and target, the on-station time, and the expected amount of drugs detected. For most of the cases we examined, the search plan generated by the pathfinding algorithm produced searcher positions close to the expected location of the target. This did not occur in some cases where the pathfinding track significantly deviated from the actual track; however, in some cases where these significant deviations occurred, the search plan produced would still be operationally effective.

The pathfinding search plans performed reasonably well for the on-station-time metric. On average, the search times produced were an hour less than equivalent plans generated using the actual intelligence data. Similarly, when we evaluate the expected amount of drugs detected, the pathfinding algorithm detects only 75 percent as much drugs as the plans using the actual intelligence. Both of these measures show degradation in performance when using the pathfinding plans; however, this illustrates that if we had to rely only on the probability heat maps for our input, we could still generate reasonable search plans using the pathfinding algorithm.

B. FUTURE WORK

There is still room to improve and build on the work in this thesis. We list several possible areas of future research below:

1. Improve the way that a user enters intelligence data into the tool. Improvements include adding drop-down menus for users to select assets, targets, or locations.
2. Automate the process of inputting data into Google Earth so that all of the data is automatically built once the model is run. Currently, the data is manually loaded into Google Earth after the model is run.
3. Develop methods to handle the “corner-cutting” effect, where estimated waypoints are off from the actual waypoints. This may involve a simple correction factor, or perhaps a fresh approach that more precisely models the movement of the blob around a corner.

4. Improve the current MARS algorithm for cases where a fast target is only on the water for a short duration. It appears that MARS defaults to OLS estimate if there are not enough data points. Perhaps we could perform interpolation to artificially inflate the sample size. Other non-MARS approaches could also be considered.
5. Formulate a conversion algorithm for cases where the probability blob never reaches 100 percent on the water. This occurs when the distance covered is small or there is significant temporal uncertainty.
6. Incorporate more sophisticated versions of the optimization algorithm into the tool. Currently, the optimization algorithm within the tool only generates one search plan for one planning cycle. Enhanced versions also incorporate interditors and multicycle planning situations. Including the multicycle plans would tremendously improve the overall performance of helicopters since they have a flight endurance of only four hours. The current optimization framework is set up more for platforms like P-3s, with long endurances and long rest periods. Helicopters have the capability to refuel while at sea and can conduct continuous flight operations at sea by swapping out flight crews.

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF REFERENCES

- Allen, C. H. (1989). Doctrine of hot pursuit: A functional interpretation adaptable to emerging maritime law enforcement technologies and practices. *Ocean Development and International Law*, 20(4), 309–341.
- Chao, I., Golden, B., & Wasil, E. (1996). Theory and methodology: The team orienteering problem. *European Journal of Operational Research*, 88, 464–474.
- Coates, N. E. (2005, March 5). The United Nations convention on the law of the sea, the United States, and international relations. Paper presented at the annual meeting of the International Studies Association, Hilton Hawaiian Village, Honolulu. Retrieved from http://www.allacademic.com/meta/p70299_index.html
- Defense Writers Group. (2010, April 10). General Douglas Fraser, United States Southern Command [Interview Transcript]. *Air Force Mag*. Retrieved from <http://www.airforcemag.com/DWG/Documents/2010/April%202010/042710FraserD.pdf>
- Faraway, J. J. (2006). *Extending the linear model with r: Generalized linear, mixed effects and nonparametric regression models*. Boca Raton, FL: Chapman & Hall/CR.
- Freedberg, S. J., Jr. (2013, May). Sequester means \$1 billion more of cocaine floods into U.S. Retrieved from <http://breakingdefense.com/2013/05/sequester-cuts-will-let-1-million-more-of-cocaine-into-us-coast-guard/>
- Friedman, J. H. (1991). Multivariate adaptive regression splines. *The Annals of Statistics*, 19(1), 1–141.
- Hansen, J., Jacobs, G., Hsu, I., Dykes, J., Dastugue, J., Allard, R., . . . Mittu, R. (2011, May). Information domination: Dynamically coupling METOC and INTEL for improved guidance for piracy interdiction. *NRL Review 2011*, 88–98.
- Hastie, T., & Tibshirani, R., & Friedman, J. (2009). *The elements of statistical learning: Data mining, inference, and prediction* (2nd ed.). Hoboken, NJ: Springer.
- Insulza, J. M. (2012). The drug problem in the Americas: Studies. The economics of drug trafficking. Organization of American States, Washington, D.C. Retrieved from http://www.cicad.oas.org/drogas/elinforme/informeDrogas2013/laEconomicaNarcotrafico_ENG.pdf
- Johnston, D. L. (1995). *Operational evaluation of surveillance effectiveness for airborne search of maritime regions* (Master's thesis). Monterey, CA: Naval Postgraduate School.

- Kelly, J. F. (2013). *Posture statement of General John F. Kelly, United States Marine Corps, Commander, United States Southern Command before the 113th Congress Senate Armed Services Committee*. Retrieved from <http://www.southcom.mil/newsroom/documents/southcom%202013%20posture%20statement%20final%20sasc.pdf>.
- Larter, D. (2014, July 26). Retiring frigates may leave some missions unfilled. *Navy Times*. Retrieved from <http://www.navytimes.com/article/20140726/NEWS/307260034>
- Mooshegian, M. (2013). *A probabilistic model of illegal drug trafficking operations in the eastern Pacific and Caribbean Sea* (Master's thesis). Monterey, CA: Naval Postgraduate School.
- National Drug Intelligence Center. (2011). National drug threat assessment 2011. Washington, DC: U.S. Department of Justice.
- Nunn, L. H. (1981). An introduction to the literature of search theory. Professional paper 305. Alexandria, VA: Center for Naval Analysis.
- Office of National Drug Control Policy. (2013). Transit zone operations. Retrieved from <http://www.whitehouse.gov/ondcp/transit-zone-operations>
- Office of National Drug Control Policy. (2013). *National southwest border counter-narcotics strategy 2013*. Washington, DC: Office of National Drug Control Policy. Retrieved from http://www.whitehouse.gov/sites/default/files/ondcp/policy-and-research/southwest_border_strategy_2013.pdf
- Pfeiff, D. M. (2009). *Optimizing employment of search platforms to counter self-propelled semi-submersibles* (Master's thesis). Monterey, CA: Naval Postgraduate School.
- Pietz, J., & Royset, J. (2013). Generalized orienteering problem with resource dependent rewards. *Naval Research Logistics*, 60, 294–312.
- Royset, J. O., & Reber, D. (2009). Optimized routing of unmanned aerial systems for the interdiction of improvised explosive devices. *Military Operations Research*, 14(4), 5–19.
- Seelke, C. R. (2010). Latin America and the Caribbean: Illicit drug trafficking and U.S. counterdrug programs. Retrieved from <http://fpc.state.gov/documents/organization/142364.pdf>

- Sidoti, D., Fernando Martinez Ayala, D., Han, X., Mishra, M., Sankavaran, S., Woosun, A., Krishna, R.P., & Klienman, D.L. (2013). Evaluating the value of information in the presence of high uncertainty. 18th ICCRTS: C2 in Underdeveloped, Degraded and Denied Operational Environments. Storrs, CT: University of Connecticut.
- United States Coast Guard. (2012). *Operations*. Coast Guard Publication 3-0. Washington, DC: Author.
- United States Naval War College. (2011, July). *Joint military operations reference guide: Forces/capabilities handbook*. Newport, RI: United States Naval War College.
- U. S. Department of the Navy. (2007). *The commander's handbook on the law of naval operations*. Naval Warfare Publication 1-14M. Norfolk, VA: Department of the Navy.
- Vansteenwegen, P., Souffriau, W., Vanden Berghe, G., & Van Oudheusden, D. (2009). A guided local search metaheuristic for the team orienteering problem. *European Journal of Operational Research*, 196(1), 118–127.
- Washburn, A. R. (2002). *Search and detection. Institute for Operations Research and the Management Sciences* (4th ed.). Cantonsville, MD: Institute for Operations Research and the Management Sciences.

THIS PAGE INTENTIONALLY LEFT BLANK

INITIAL DISTRIBUTION LIST

1. Defense Technical Information Center
Ft. Belvoir, Virginia
2. Dudley Knox Library
Naval Postgraduate School
Monterey, California